SIGGRAPH ASIA 2009

革新の波動

the pulse of innovation

Amortized Supersampling

Lei Yang [H], Diego Nehab [M], Pedro V. Sander [H],
Pitchaya Sitthi-amorn [V], Jason Lawrence [V],
Hugues Hoppe [M]

[H] 香港科技大學 THE HONG KONG UNIVERSITY OF SCIENCE AND TECHNOLOGY
[M] Microsoft Research
[V] University of Virginia

SIGGRAPH ASIA 2009

Dec. 18, 2009, Pacifico Yokohama, Japan

Outline

- Problem
- Amortized supersampling – basic approach
- Challenge - the resampling blur
- Our algorithm
- Results and conclusion

Problem

- Shading signals not band-limited
  - Procedural materials
  - Complex shading functions
- Band-limited version (analytically antialiased)
  - Ad-hoc
  - Difficult to obtain

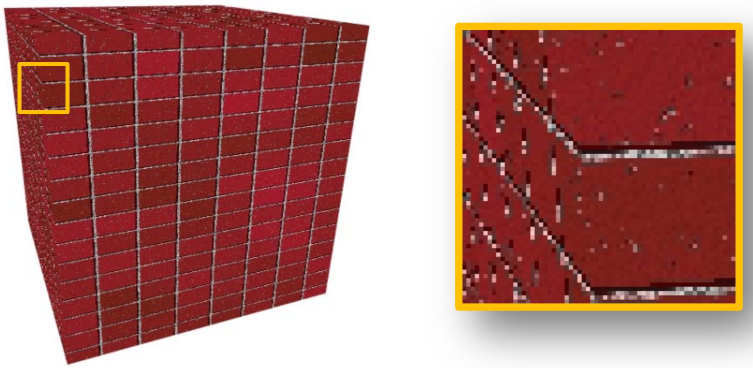

Problem

- Supersampling
  - General antialiasing solution
  - Compute a Monte-Carlo estimator

$$f_N[p] \leftarrow \frac{1}{N} \sum_{i=1}^{N} s_t[p]$$

  - Can be prohibitively expensive

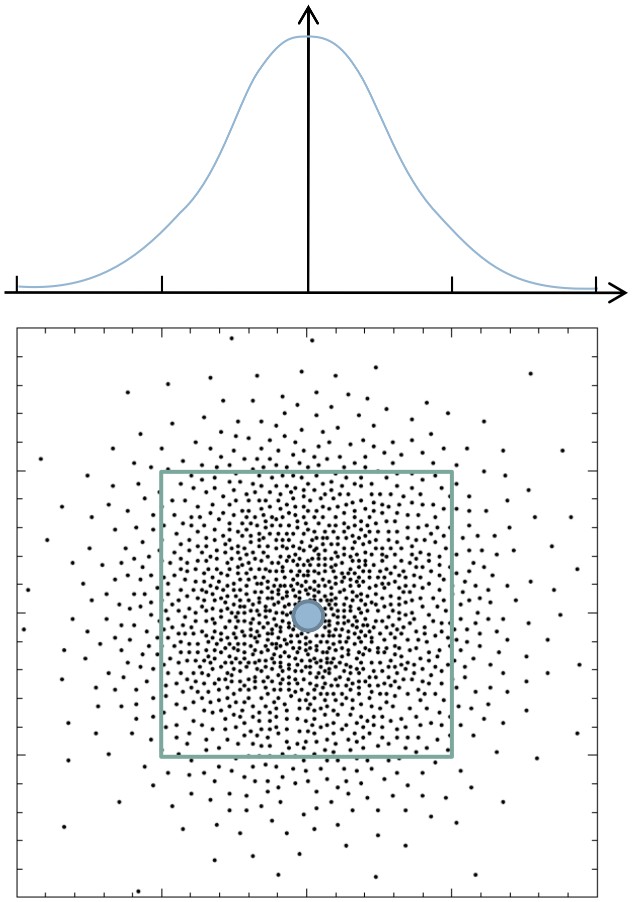

Accelerating Supersampling

- Shading functions usually vary slowly over time
- Reuse samples from previous frames
  - Reprojection
  - Generate only one sample every frame

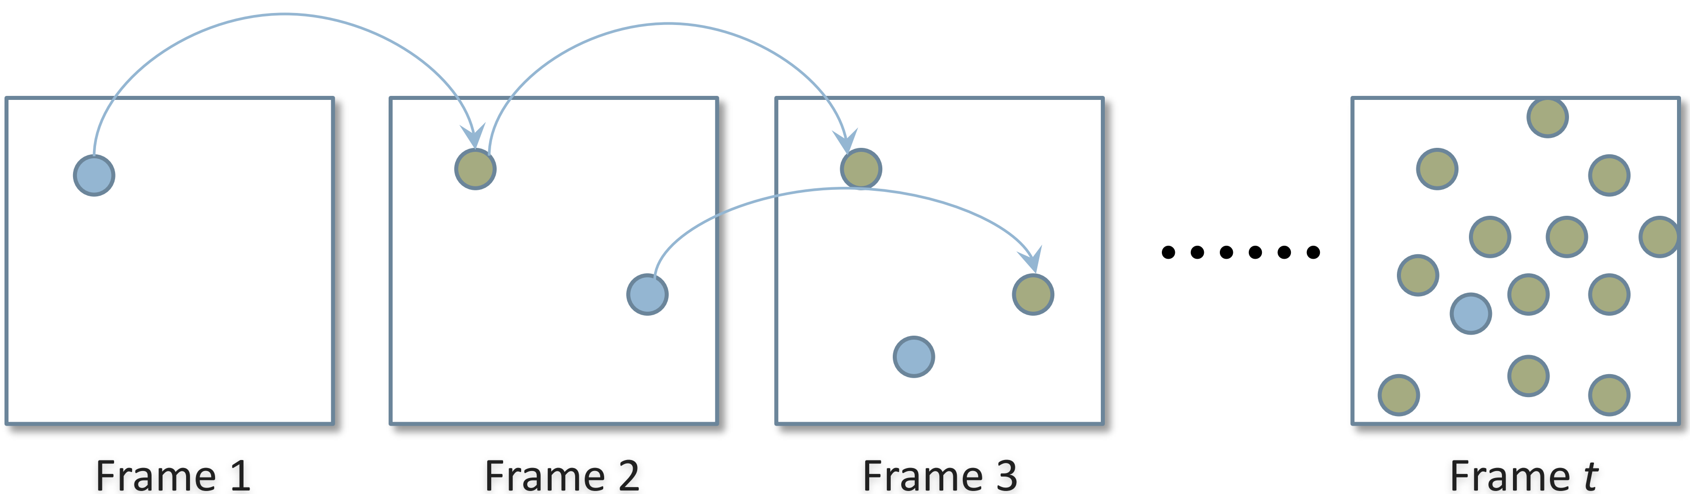

Amortized Supersampling

- Cannot afford to store all the samples from history
- Keep only a running tally f_t per pixel
  - Update it every frame using exponential smoothing

$$f_t[p] \leftarrow (\alpha)\boxed{s_t[p]} + (1-\alpha)\boxed{f_{t\text{-}1}\left(\pi_{t\text{-}1}(p)\right)}$$

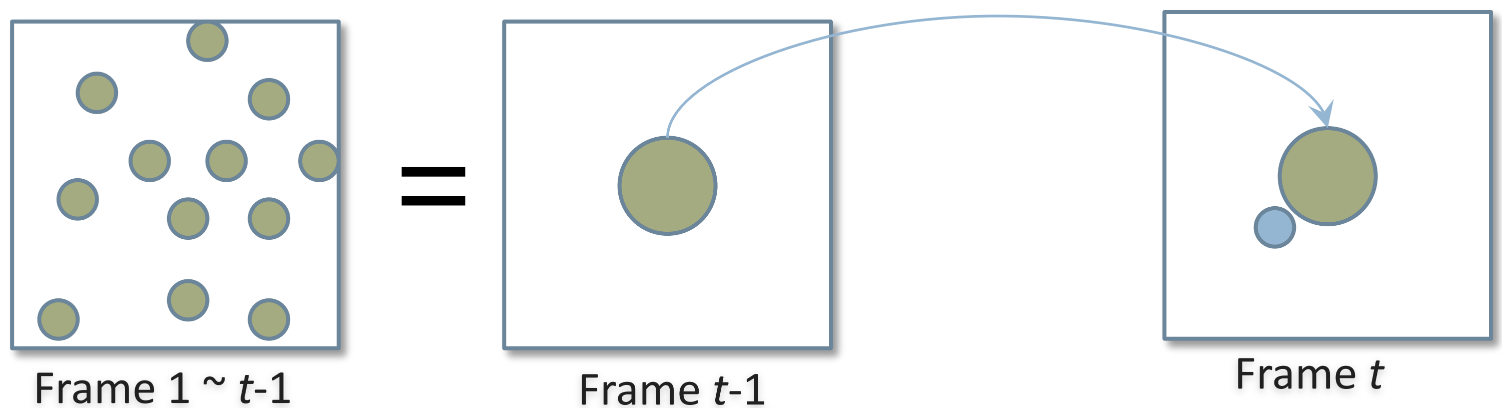

Reverse Reprojection [Nehab07, Scherzer07]

- Compute previous location $\pi_{t\text{-}1}(p)$ of point p
- A bilinear texture fetch for the previous value $f_{t\text{-}1}\left(\pi_{t\text{-}1}(p)\right)$
  - Check depth for occlusion changes

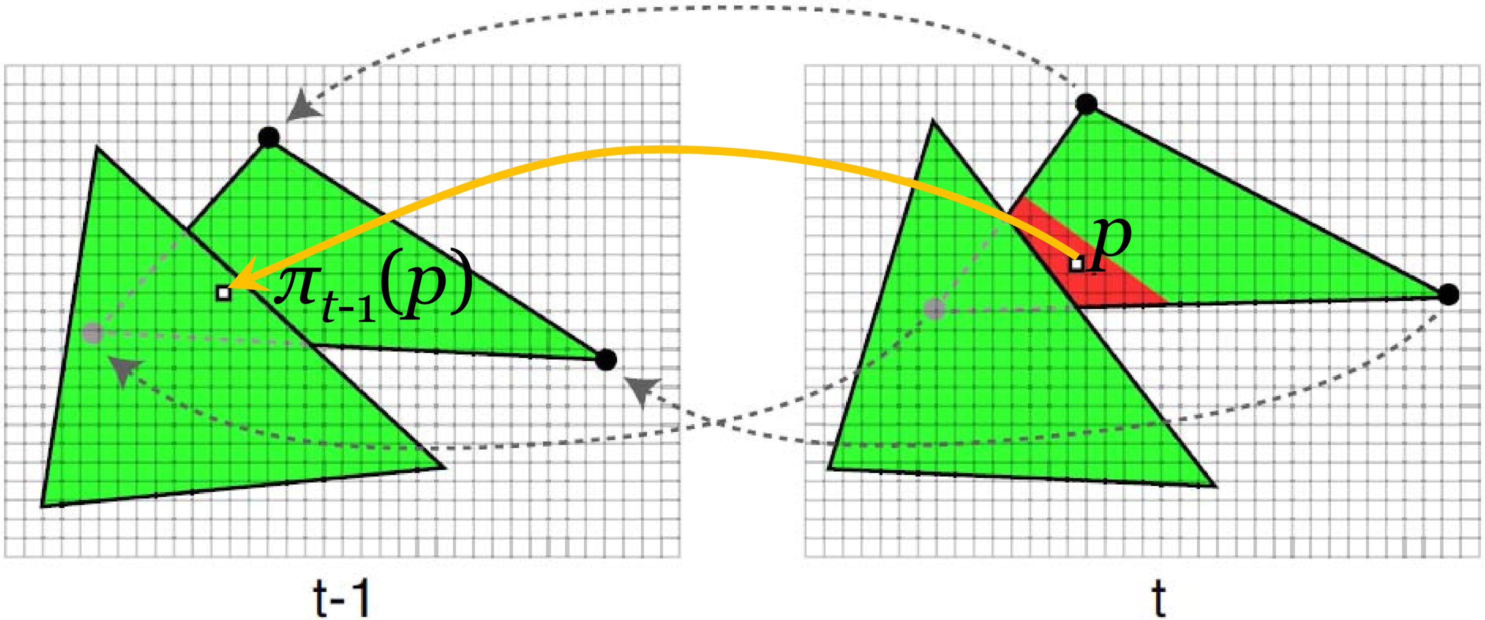

Effect of the smoothing factor α

- Larger α: less history, more aliasing/noise
- Smaller α: more history, less aliasing/noise
- Equal weight of samples: $\alpha = \frac{1}{t}$

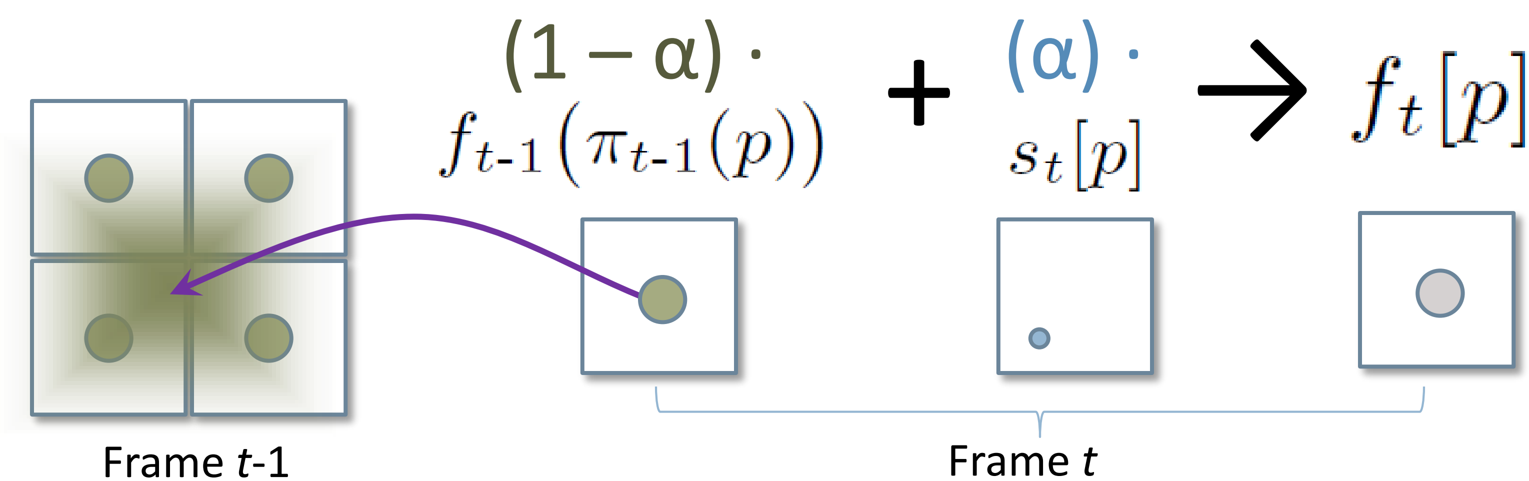

An artifact of recursive reprojection

- Blur due to repeated bilinear interpolation

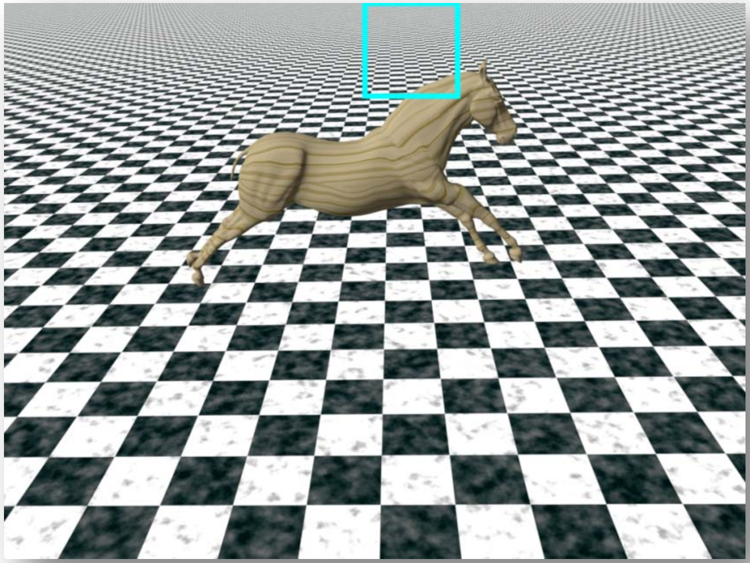

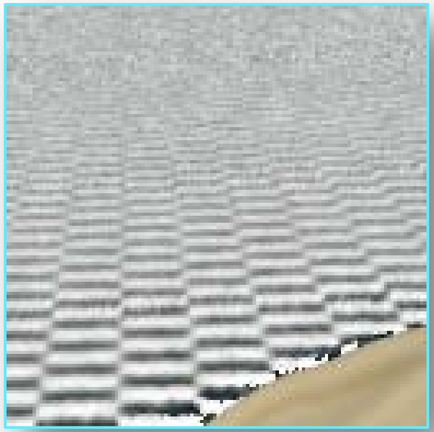

SS through Reprojection

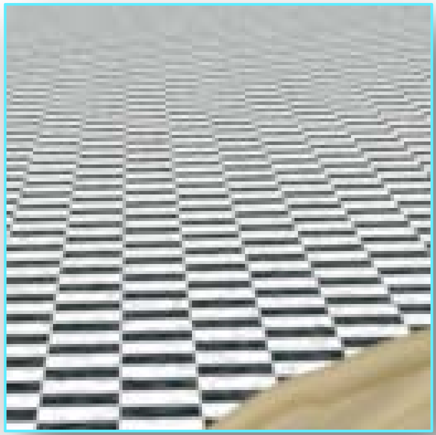

Ground truth

Factors of the blur

- Fractional pixel velocity $v = (v_x, v_y)$

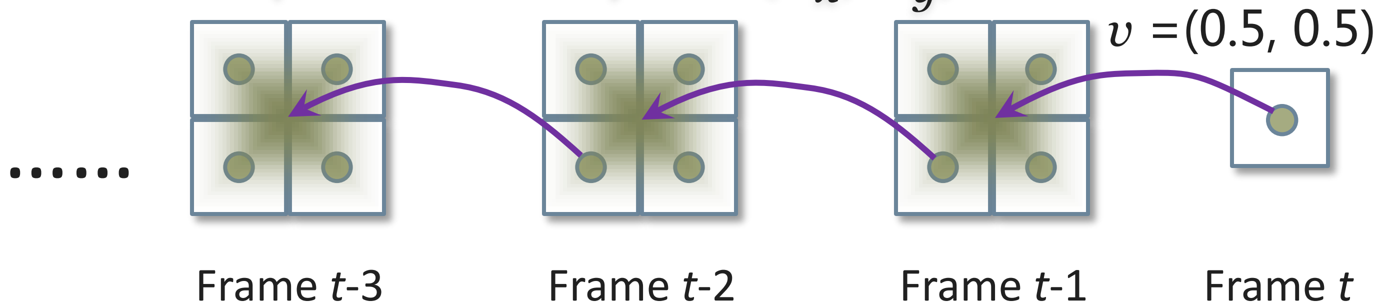

- Exponential smoothing factor α

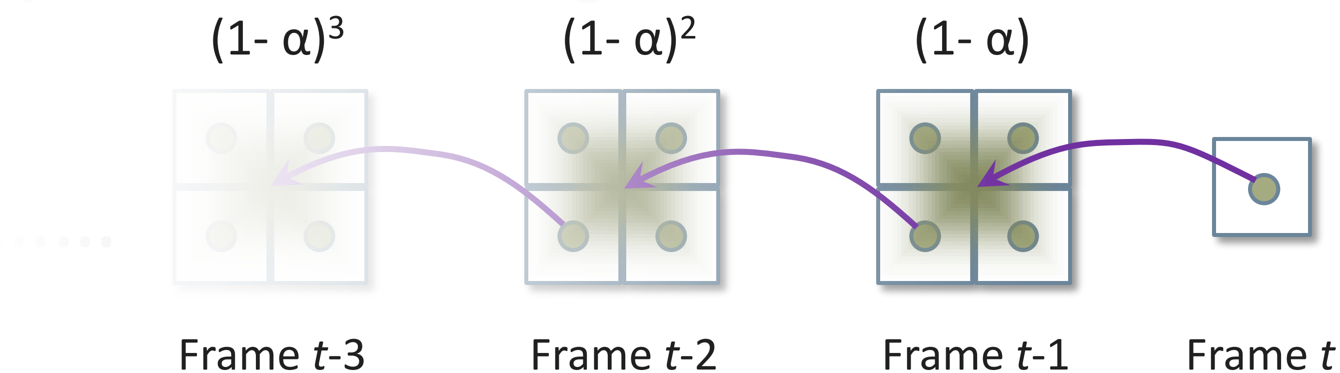

The amount of blur

- The expected blur variance is (derivation in the appendix)

$$\sigma_v^2 = \sigma_G^2 + \frac{1-\alpha}{\alpha} \frac{v_x(1-v_x) + v_y(1-v_y)}{2}$$

- Approaches for reducing the blur:
  1. Increase resolution of the history buffer
  2. Avoid bilinear resampling whenever possible
  3. Limit α when needed

(1) Increase resolution

- Option 1:
  Keep a history buffer at high resolution (2x2)
  - Have to update it every frame ☹
- Option 2:
  Keep 4 subpixel buffers at normal resolution
  - Only update one of them each frame ☺

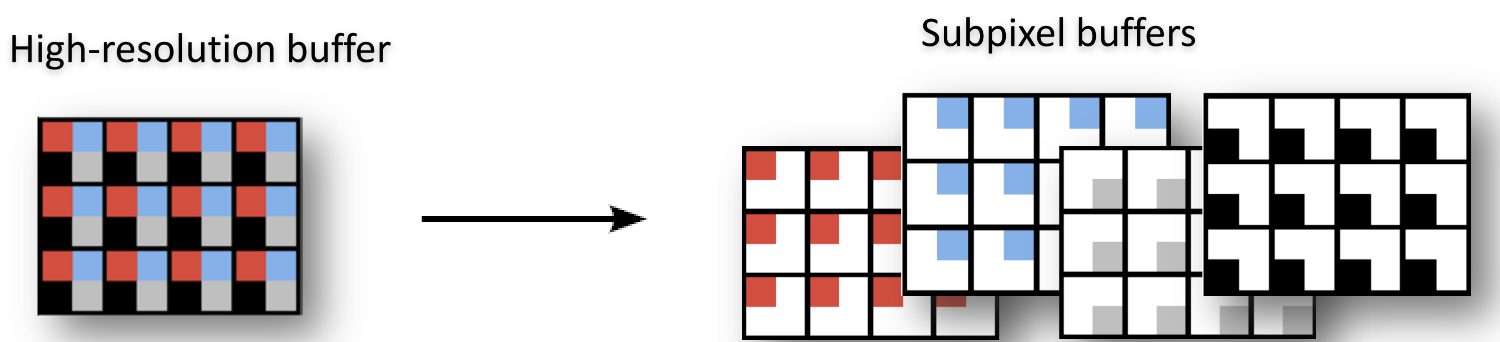

Subpixel buffers

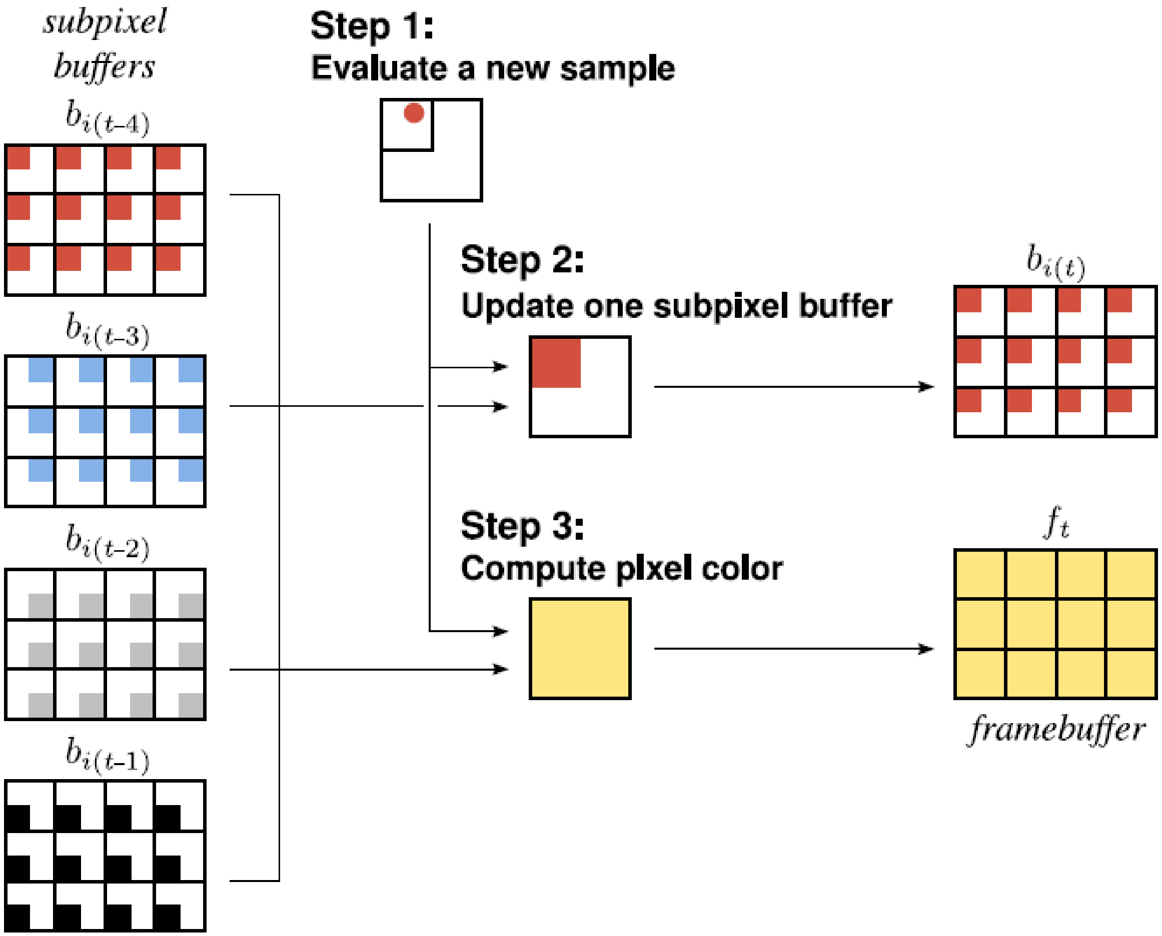

(2) Avoid bilinear sampling

- Reconstructing from subpixel buffers
  - Forward reproject the samples from 4 subpixel buffers to the current subpixel quadrant
  - Weight them using a tent function
  - GPU approximation/acceleration

Reconstruction scheme

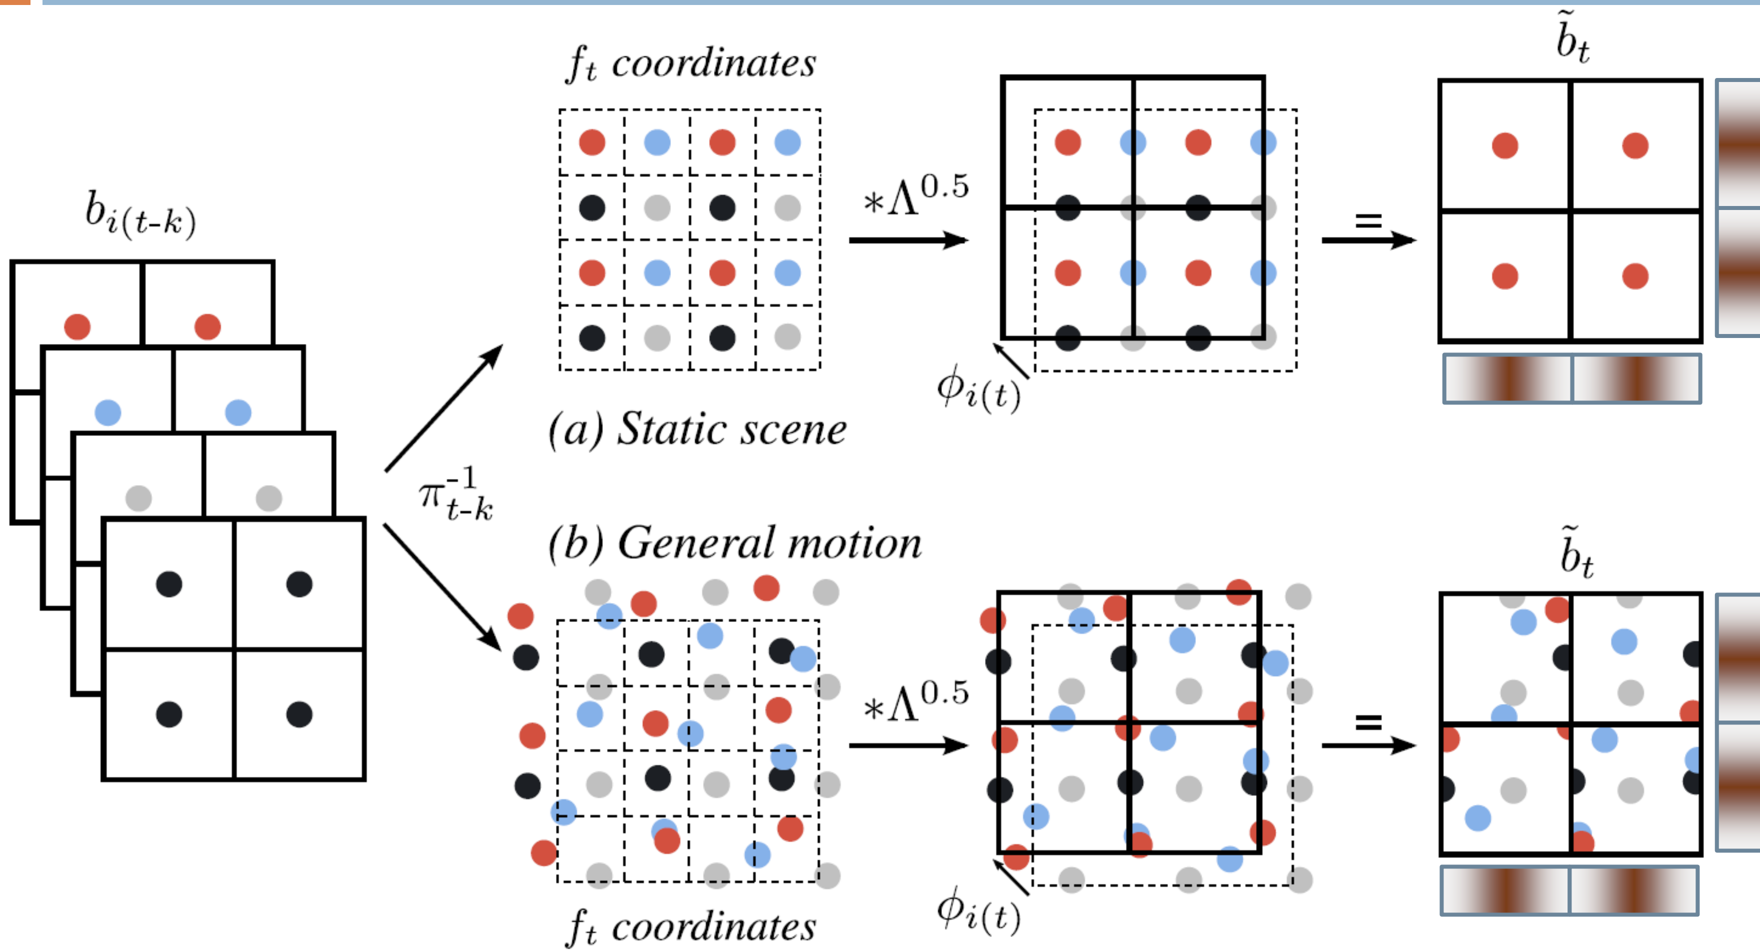

(3) Limiting blur via bounding α

- Derive a relationship between
  - Blur variance σ^2
  - Motion velocity v and α
- Analytic relationship is not attainable
  - Numerical simulation and tabulate
- Bound α for limiting σ^2 no larger than τ_b

Tradeoff of blur and aliasing

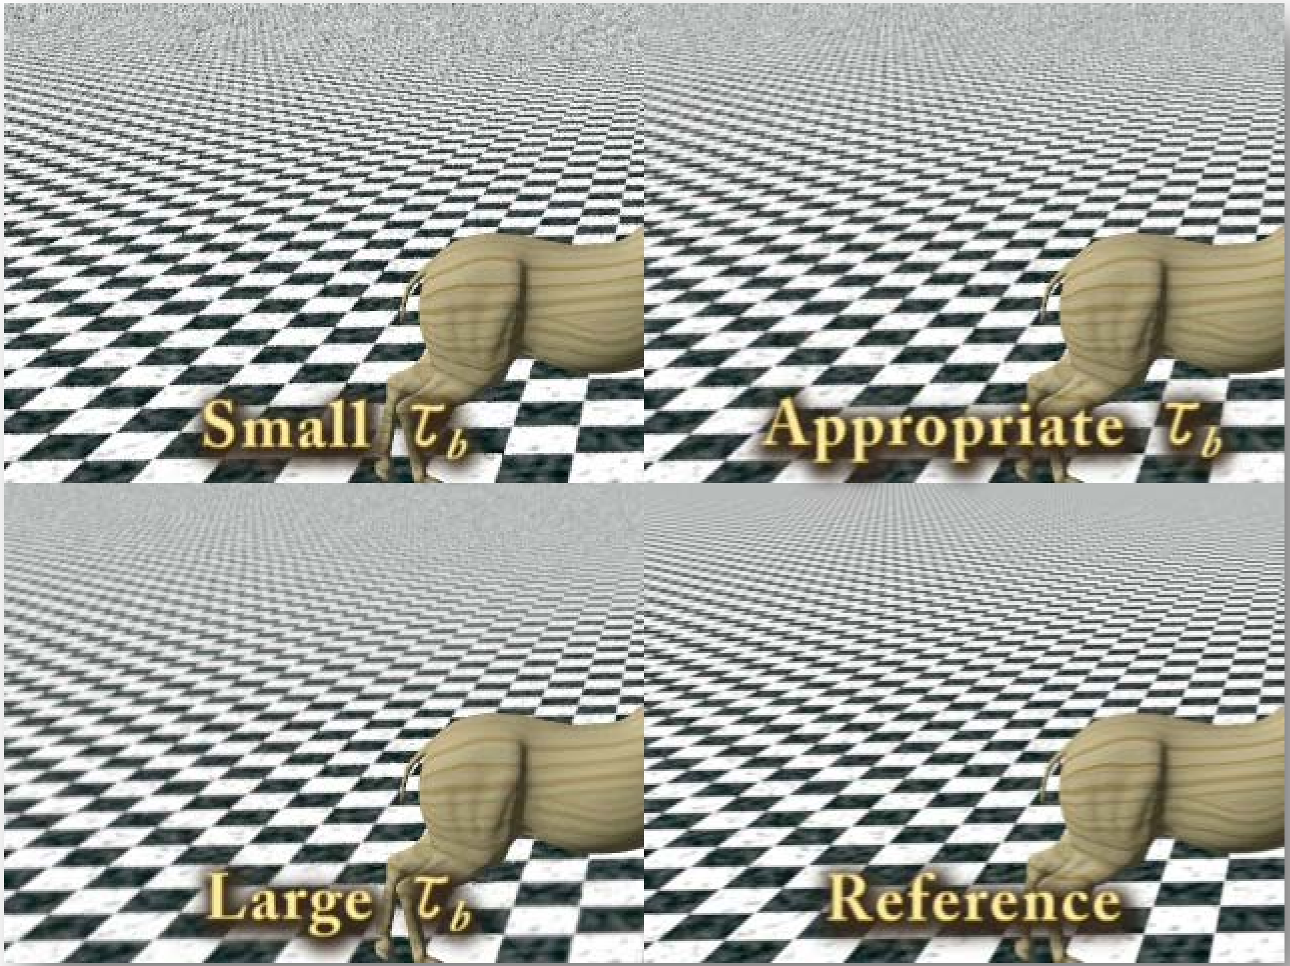

Adaptive evaluation

- Newly disoccluded pixels are prone to aliasing
- Additional shading for subpixels that fail in reconstruction

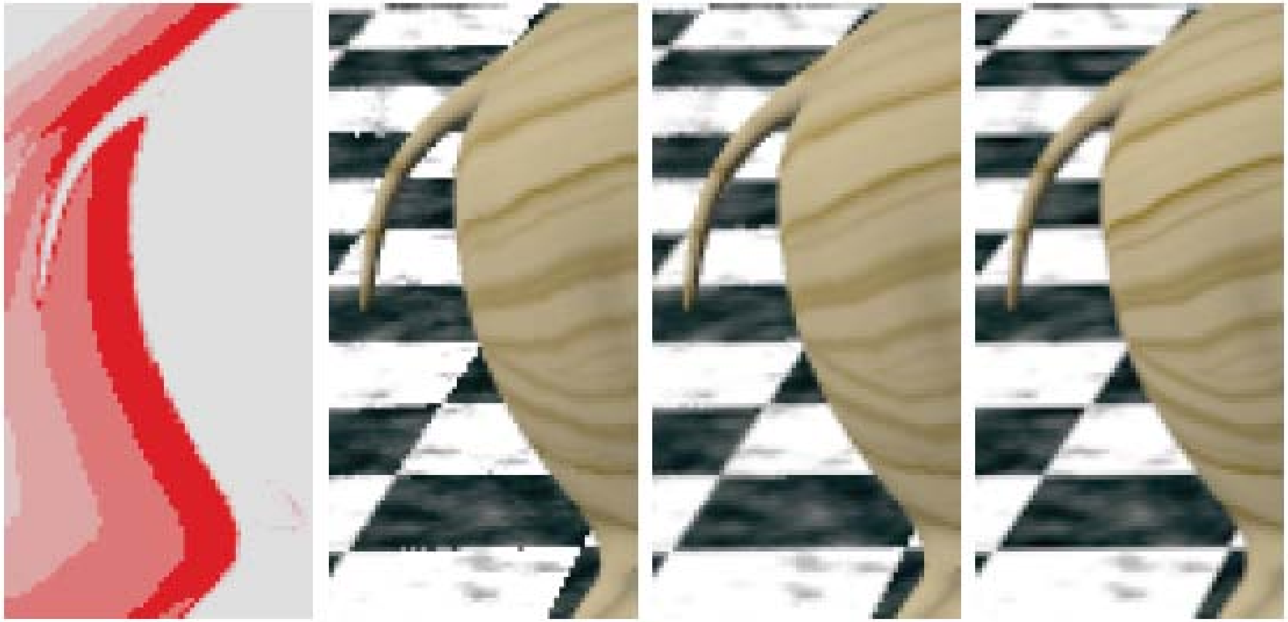

(a) Eval map *(b) One eval* *(c) Adaptive eval* *(d) Reference*

Accounting for signal changes

- Detect fast signal change
  - React by more aggressive update
- Estimate residual ε between:
  - Current sample s_t (aliased/noisy)
  - History estimate f_t
- Blur the residual estimate to remove aliasing/noise
- Bound α for limiting ε no larger than τ_ε

Tradeoff of signal lag and aliasing

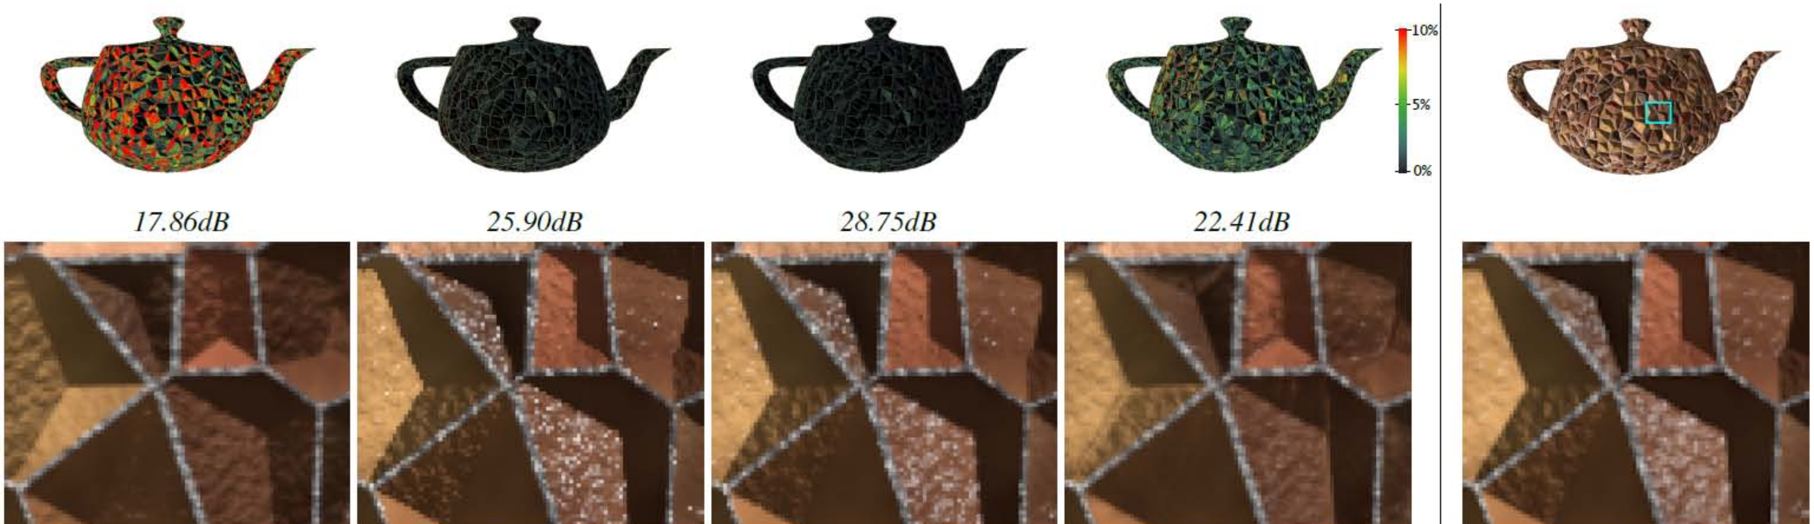

(a) No signal adaptation (b) Small τ_ϵ (c) Good τ_ϵ (d) Large τ_ϵ (e) Reference

Results

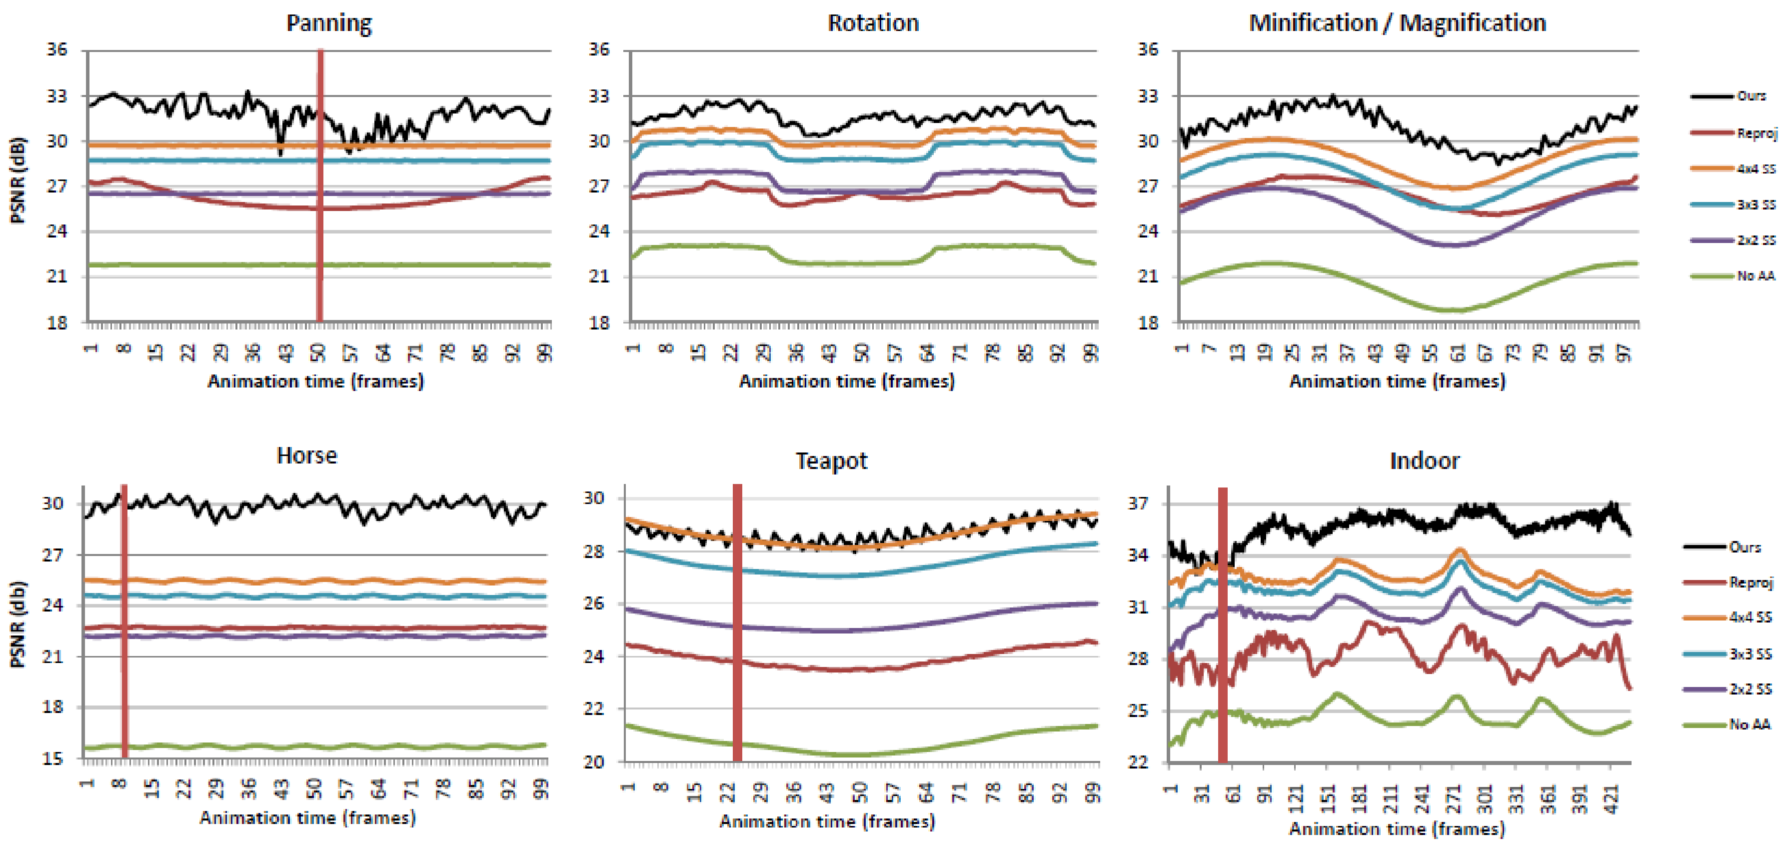
Panning
Rotation
Minification / Magnification
Horse
Teapot
Indoor
PSNR (dB)
PSNR (db)
Animation time (frames)
Ours
Reproj
4x4 SS
3x3 SS
2x2 SS
No AA

Results

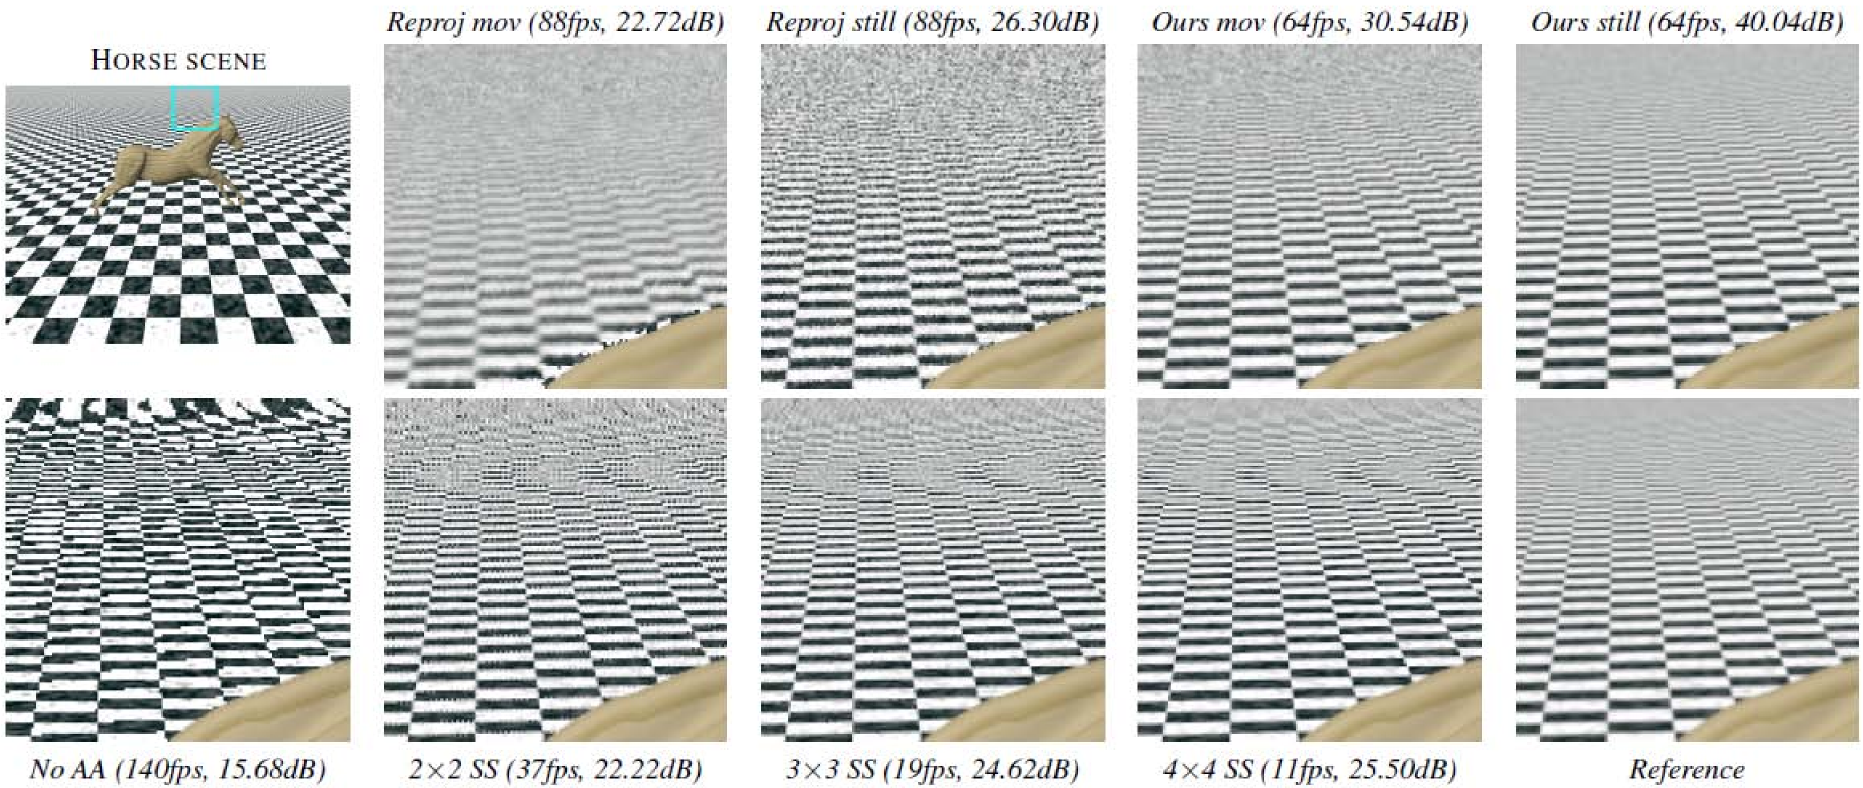
HORSE SCENE
Reproj mov (88fps, 22.72dB)
Reproj still (88fps, 26.30dB)
Ours mov (64fps, 30.54dB)
Ours still (64fps, 40.04dB)
No AA (140fps, 15.68dB)
2×2 SS (37fps, 22.22dB)
3×3 SS (19fps, 24.62dB)
4×4 SS (11fps, 25.50dB)
Reference

Results

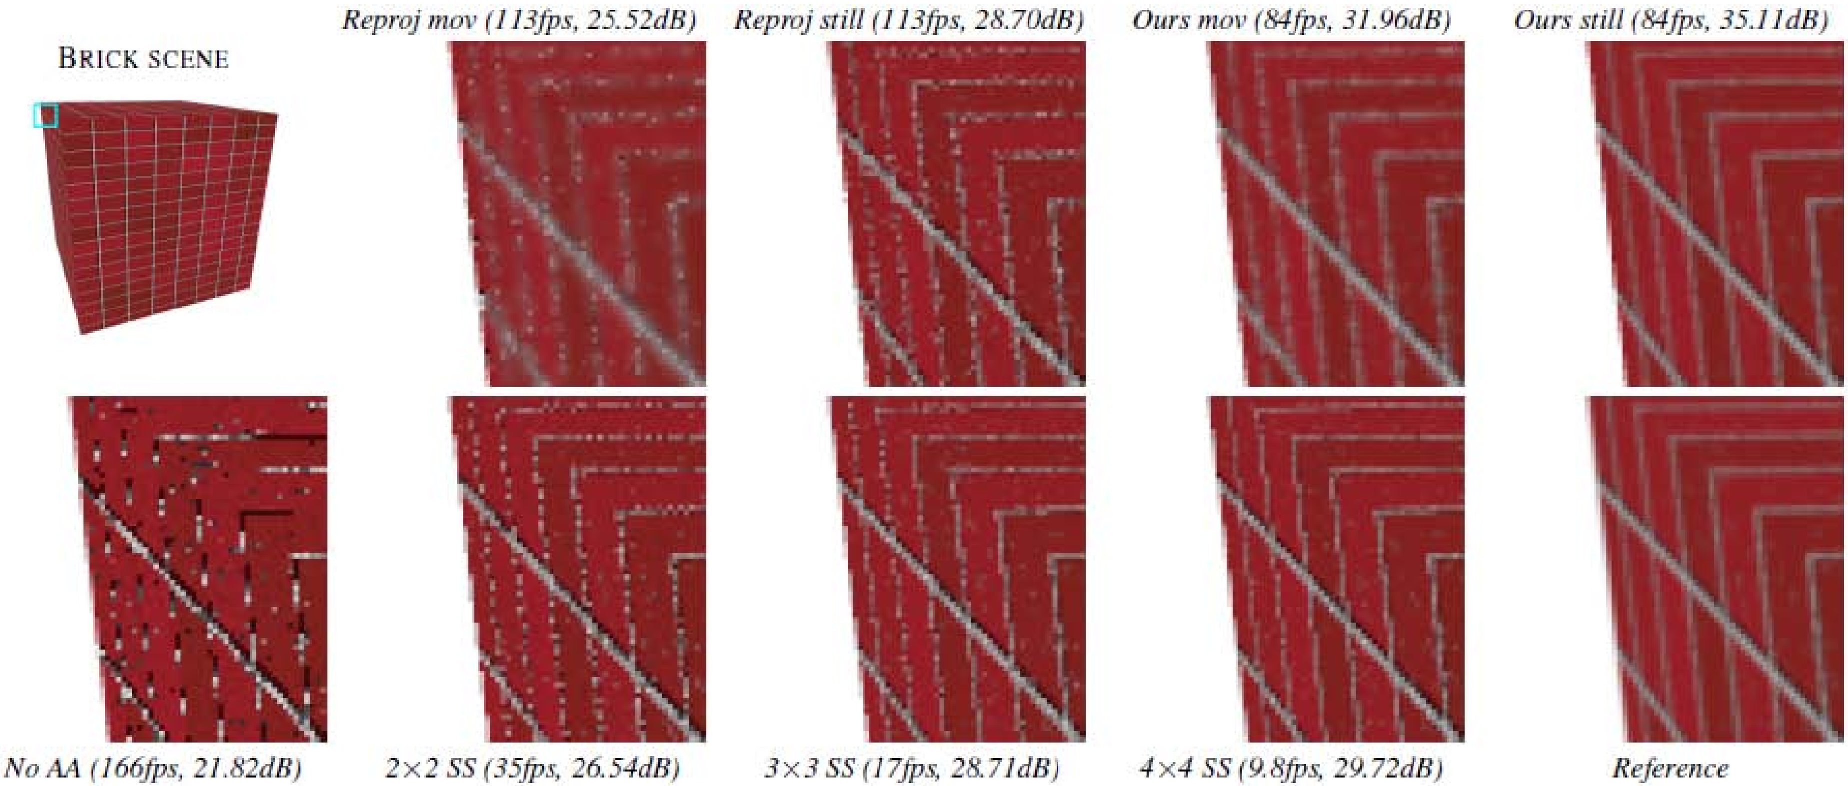

Brick scene
Reproj mov (113fps, 25.52dB)
Reproj still (113fps, 28.70dB)
Ours mov (84fps, 31.96dB)
Ours still (84fps, 35.11dB)
No AA (166fps, 21.82dB)
2×2 SS (35fps, 26.54dB)
3×3 SS (17fps, 28.71dB)
4×4 SS (9.8fps, 29.72dB)
Reference

Results

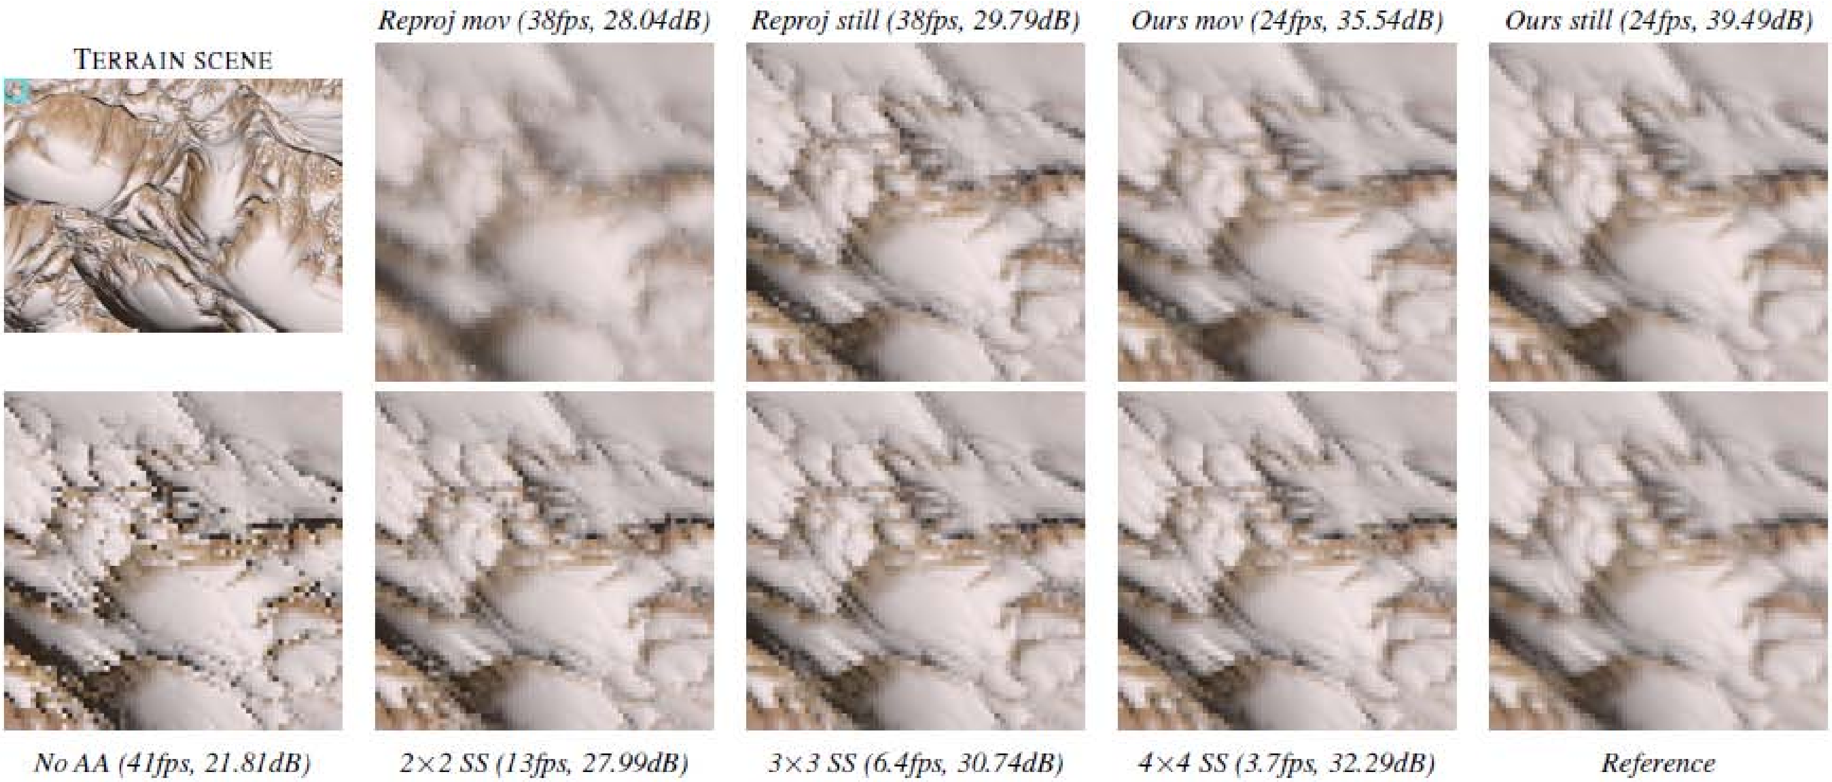
Reproj mov (38fps, 28.04dB)
Reproj still (38fps, 29.79dB)
Ours mov (24fps, 35.54dB)
Ours still (24fps, 39.49dB)
TERRAIN SCENE
No AA (41fps, 21.81dB)
2×2 SS (13fps, 27.99dB)
3×3 SS (6.4fps, 30.74dB)
4×4 SS (3.7fps, 32.29dB)
Reference

Results

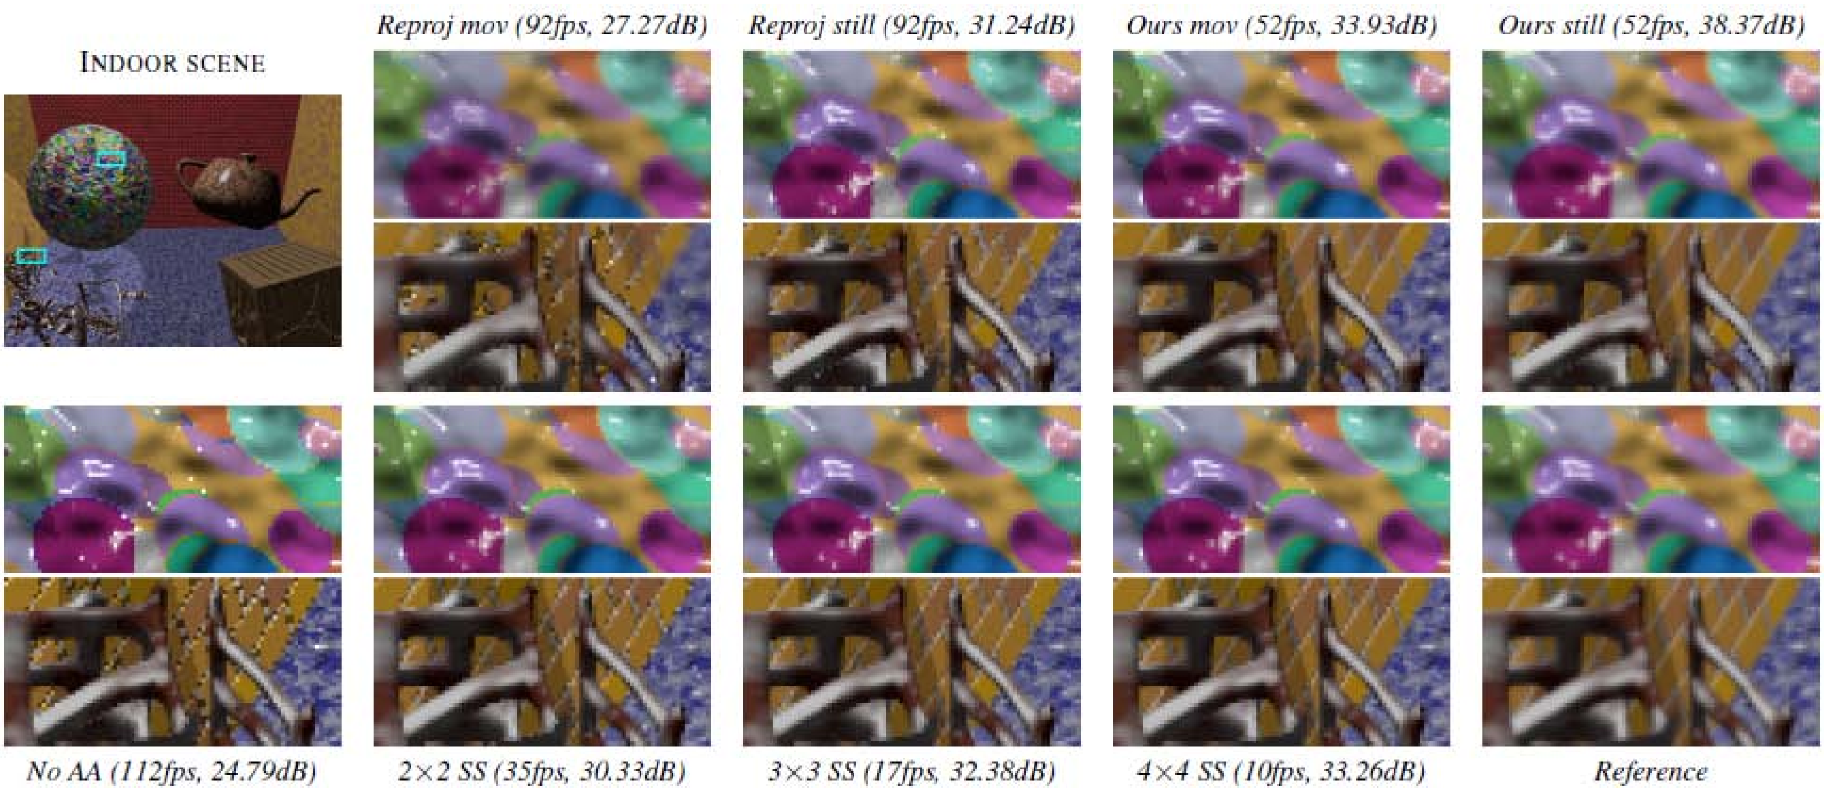
INDOOR SCENE
Reproj mov (92fps, 27.27dB)
Reproj still (92fps, 31.24dB)
Ours mov (52fps, 33.93dB)
Ours still (52fps, 38.37dB)
No AA (112fps, 24.79dB)
2×2 SS (35fps, 30.33dB)
3×3 SS (17fps, 32.38dB)
4×4 SS (10fps, 33.26dB)
Reference

Conclusion

- A real-time scheme for amortizing supersampling costs
  - Quality comparable to 4x4 stratified supersampling
  - Speed is 5x-10x of 4x4 supersampling
  - A single rendering pass
- Future work
  - A broader range of temporal effects
  - Edge AA and motion blur

Questions?

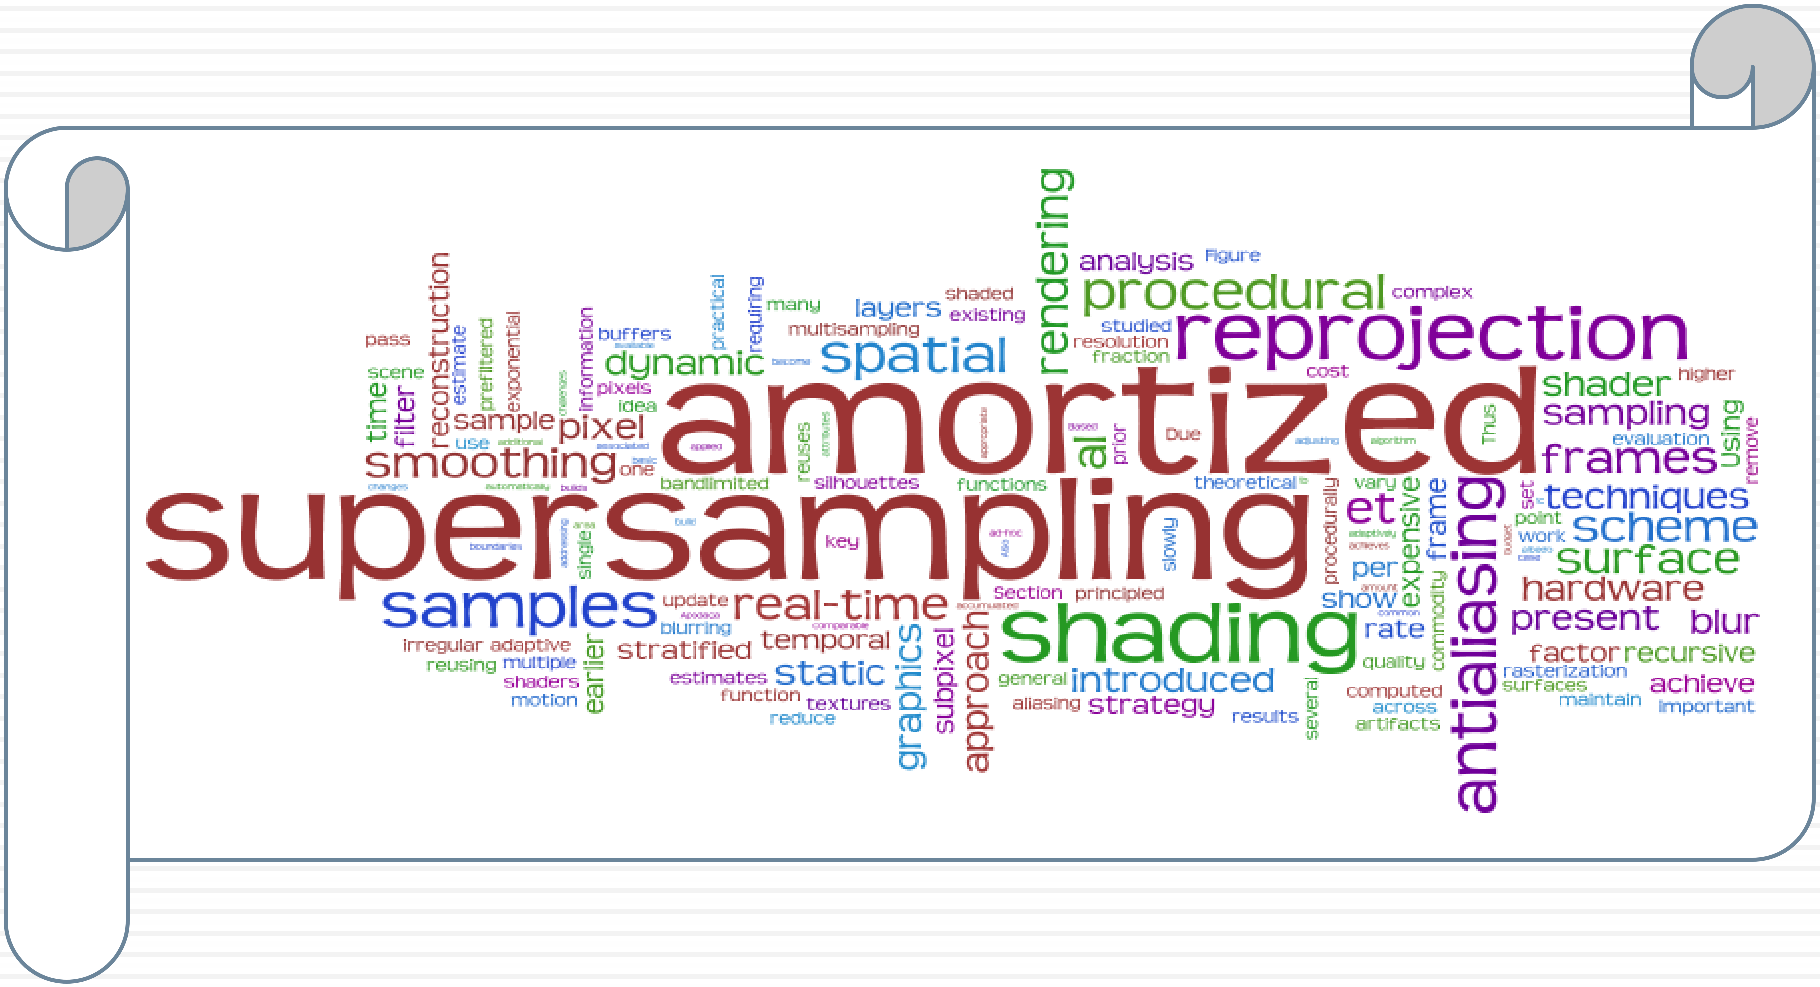
supersampling
amortized
reprojection
shading
antialiasing
procedural
spatial
rendering
samples
smoothing
real-time
dynamic
pixel
sample
frames
sampling
shader
techniques
scheme
surface
hardware
present
blur
factor
recursive
achieve
important
maintain
rasterization
surfaces
introduced
strategy
static
temporal
stratified
graphics
subpixel
approach
analysis
Figure
complex
layers
existing
shaded
many
multisampling
buffers
practical
requiring
information
reconstruction
estimate
prefiltered
exponential
pass
scene
time
filter
use
one
pixels
idea
bandlimited
silhouettes
functions
theoretical
key
reuses
al
prior
Due
cost
higher
evaluation
using
remove
studied
resolution
fraction
et
per
show
rate
quality
vary
expensive
frame
commodity
procedurally
slowly
Section
principled
update
blurring
irregular
adaptive
reusing
multiple
shaders
motion
earlier
single
estimates
function
textures
reduce
general
aliasing
results
several
computed
across
artifacts
set
point
work
Thus

GPU acceleration for reconstruction

- Forward-reproject 16 samples to the tent→
- Reverse-reproject the tent to 4 subpixel buffers
  - A single bilinear fetch in each