mental images®

Rendering Imagination Visible™

Realism in real-time will require...

June 27, 2010

Alexander Keller
Chief Scientist

www.mentalimages.com

alex@mental.com

mental images GmbH
Fasanenstrasse 81
10623 Berlin
Germany

mental images, Inc.
One Embarcadero Center, Suite 500
San Francisco, CA 94111
USA

mental images Pty. Ltd.
1/140 Bourke Street
Melbourne, Victoria 3000
Australia

Realism in real-time will require...

mental images®

Programmable memory interface and scheduler

- efficient path tracing through latency hiding
  - rays diverge unless they are parallel
    - in fact rays are created as incoherent as possible
    - transparent use of cache is not efficient
    - explicitly specify what goes through what memory, especially the cache
  - shading touches even more memory
    - layered material models with parameters from textures
    - texture caching will not work efficiently, either
    - efficient ”small” bit-width random memory access
- scalable parallelization beyond shared memory
  - path tracing not at all ”embarrassingly parallel”
  - parallel acceleration data structure construction
  - programmable collaborative operations
    - for example parallel reordering and sorting primitives

Realism in real-time will require...

Level of detail

- selection criteria other than view dependent
- axis-aligned boxes instead of micropolygons
  - overlap guarantees no cracks without any stitching effort
  - if any fixed function unit, then fixed point ray tracing

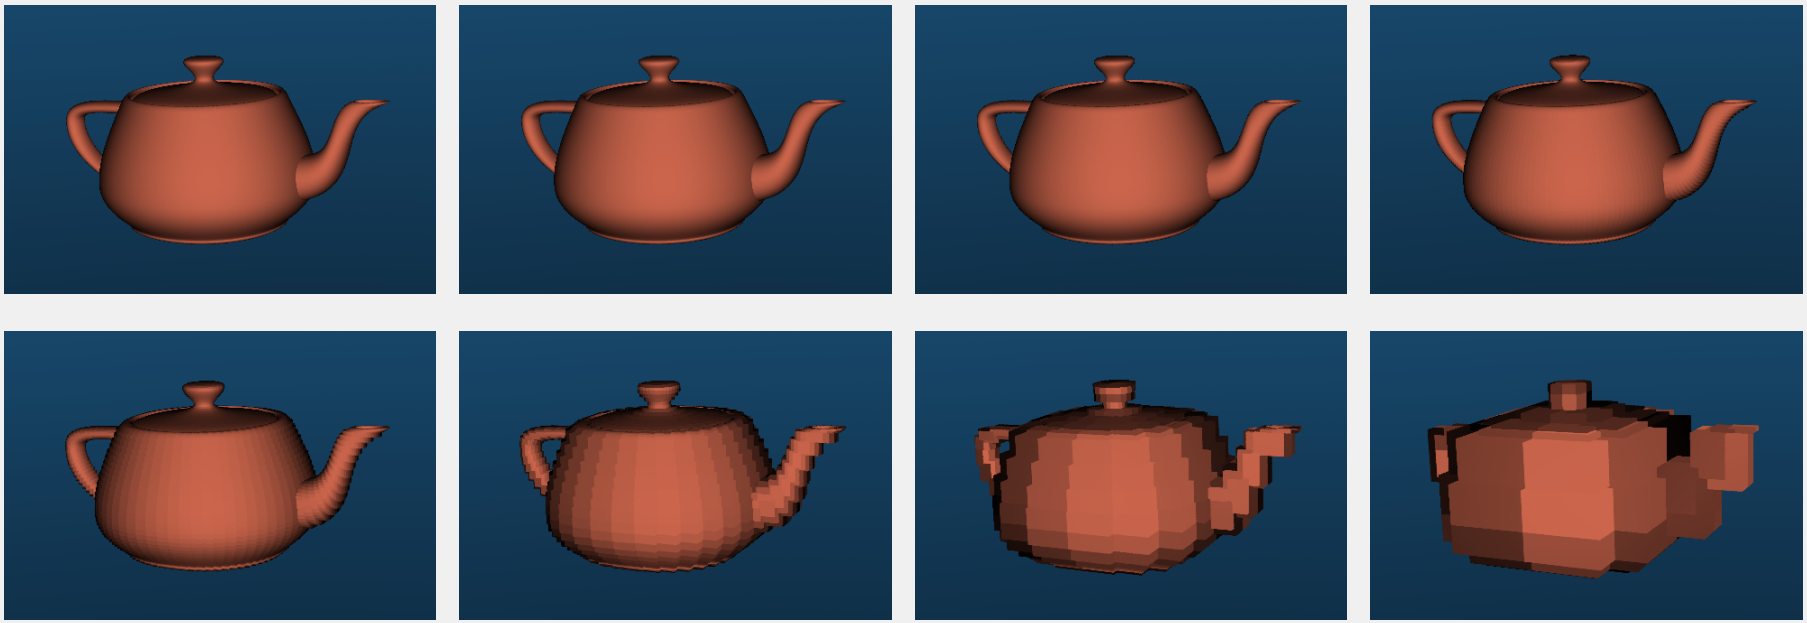

Anti-aliasing

- 1 random sample per pixel

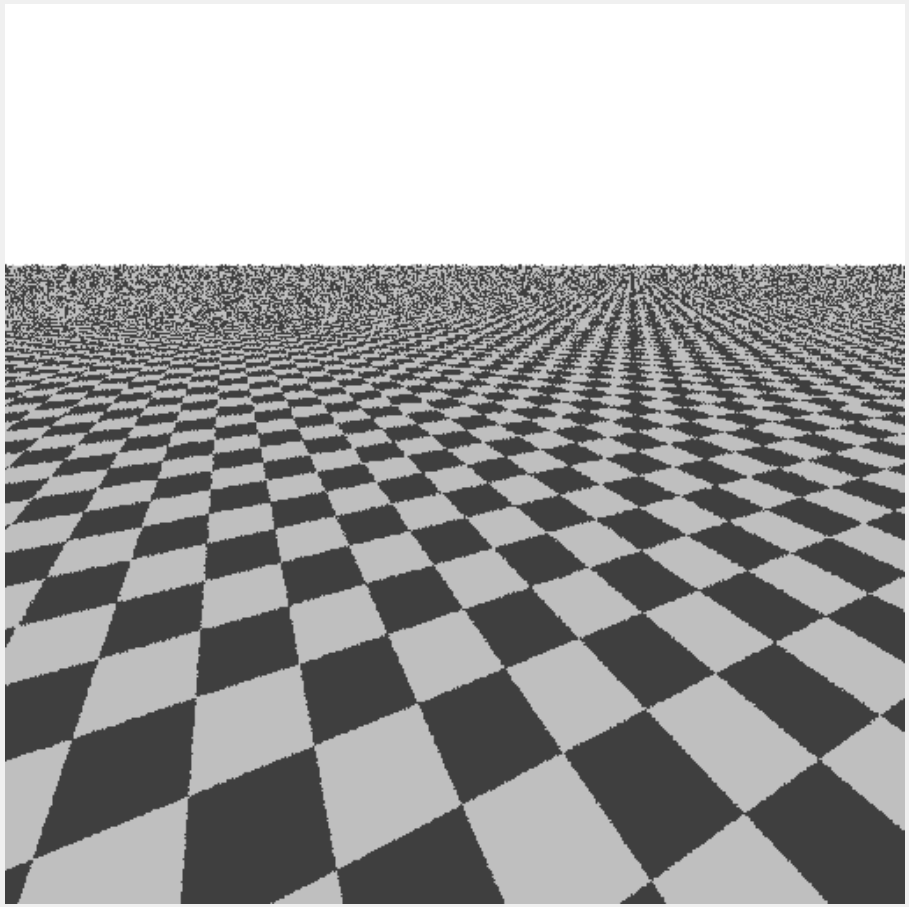

- artifacts covered by noise
- however, freckled edges

Anti-aliasing

- 16 random samples per pixel

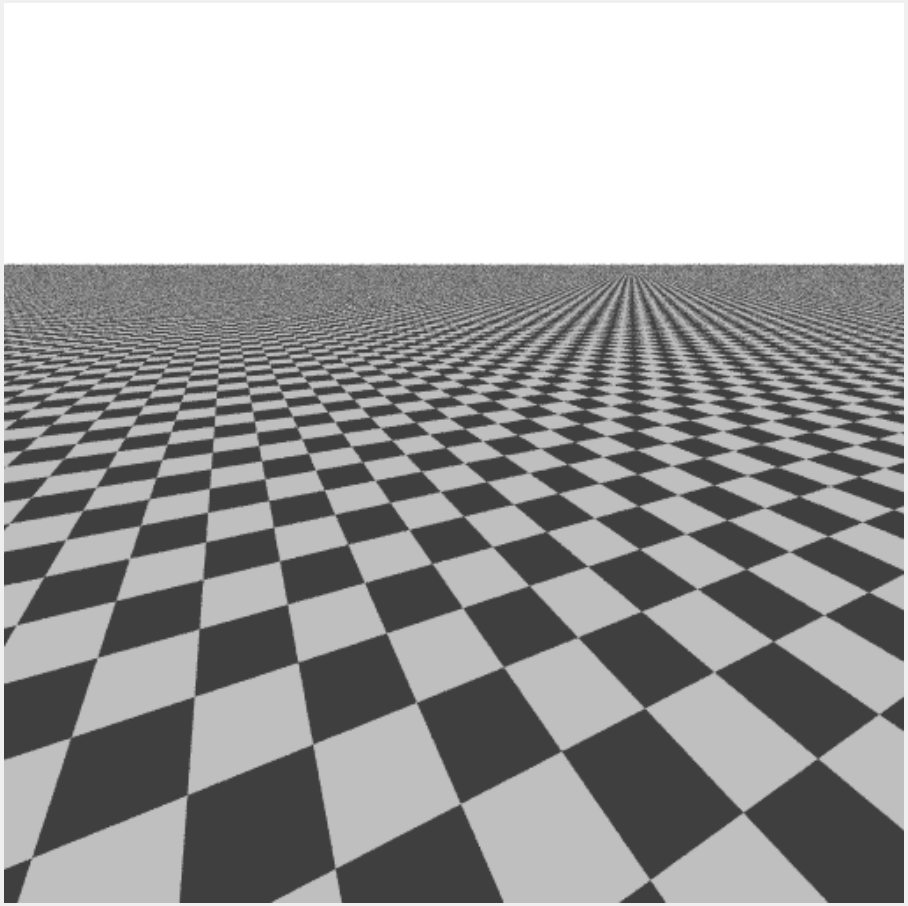

- slower
- better averages
- looks better

Anti-aliasing

- 4 × 4 stratified random samples per pixel

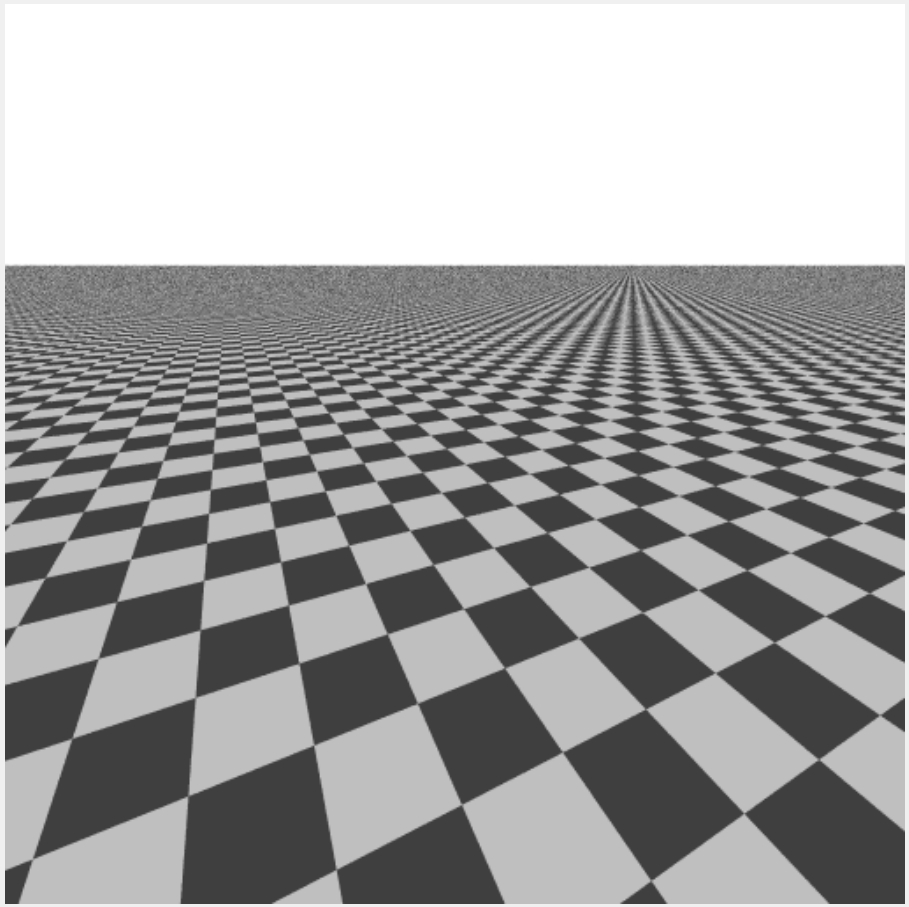

- often converges faster

Anti-aliasing

- 1024 × 1024 stratified random samples per pixel

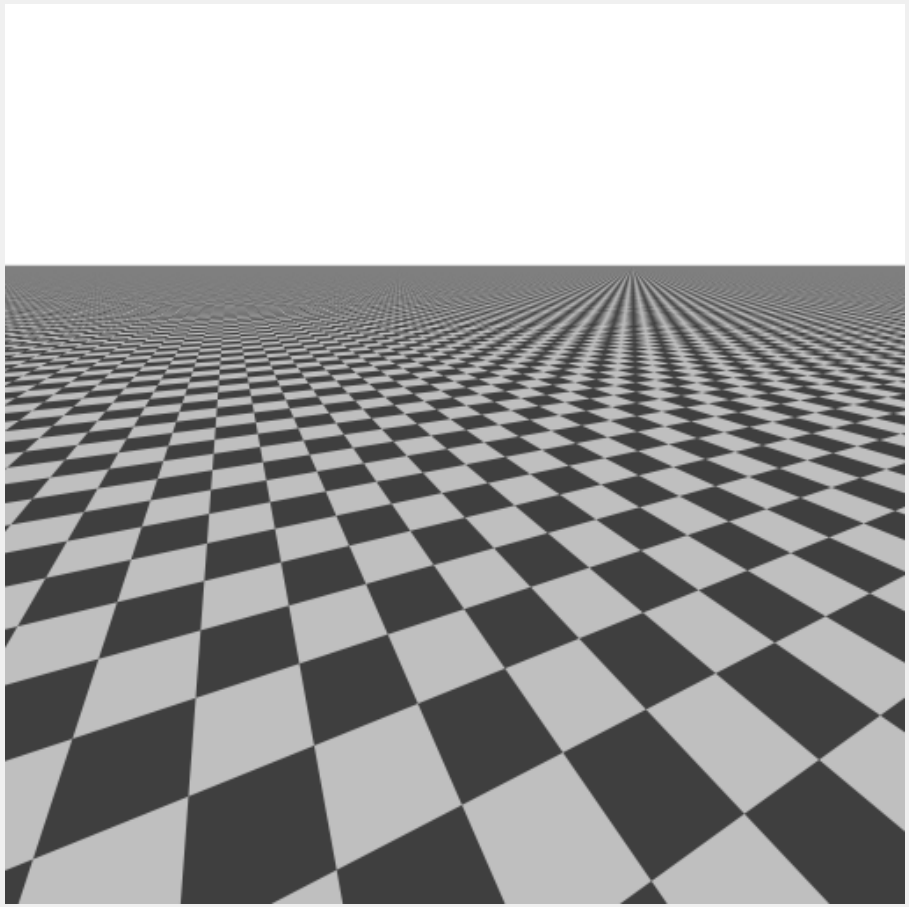

- converges faster to aliasing
- 2 × 2 pixel blocks
  - at the horizon

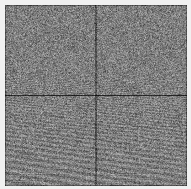

  - in the middle

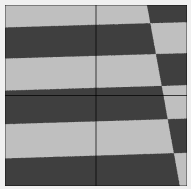

  - in the front

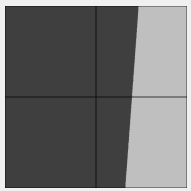

Parametric Quasi-Monte Carlo Integration

'For every randomized algorithm, there is a clever deterministic one.'
Harald Niederreiter, Claremont, 1998.

- misconceptions about random sampling
  - only unbiased algorithms are good
    - convergence speed must be sacrificed due to independence
  - only random sampling can avoid aliasing
    - the aliases will return sooner or later...
  - only blue noise samples will do the job
    - cursed by dimension
- consistent deterministic algorithms
  - faster, unconditional convergence
  - much larger class of algorithms
  - reproducible
  - simple to parallelize
  - simple to use, even without any parameters