Global Illumination IV

Photon Mapping

G. Papaioannou 2010 - MSc Graphics Course

Photon Mapping (1)

A practical and robust **two-pass** algorithm:

- Traces illumination paths **both** from the lights and from the viewpoint
- **Caches and reuses** illumination values in a scene for efficiency
- Light source flux is spread among outgoing light paths ("photon" trajectories) and carried to the scene
- The density of the photons is captured by the camera paths

- Photon mapping is **biased**

Photon Mapping (2)

- First pass:
  - Photons are traced from the light sources into the scene
  - They carry **flux** information
  - They are cached in a data structure, called the **photon map**
- Second pass:
  - an image is rendered using the information stored in the photon map

Benefits of Particle Tracing

- Does not require Lambert surfaces
- Can interact with participating media
  - Participating media (i.e. gaseous void space or heterogeneous permeable media) can attenuate particles
  - Participating media can store photons, hence they can be lit!
- Can model wavelength-dependent behaviour of light and materials statistically
- Can model specular inter-reflection effects (e.g. **caustics**)

The Photon Map

- Stores light ray-surface hits
  - Position
  - Incident direction
  - Flux
- Uses a special, separate 3D data structure, so:
  - Decouples photon storage from surface parameterization
  - Handles arbitrary geometry, including procedural geometry
  - Increases the practical utility of the algorithm
  - Not prone to meshing artifacts

k-D Tree as a Photon Map

- A k-D Tree is a binary tree that partitions the k-dimensional space according to one dimension at a time
- Each new inserted photon (hit) potentially creates a splitting plane in the k-d tree (up to a maximum level)
- The dimension of maximum cell size is chosen as a splitting plane
- The photon map is finally balanced to minimize search time

k-D Tree as a Photon Map – 2D Example

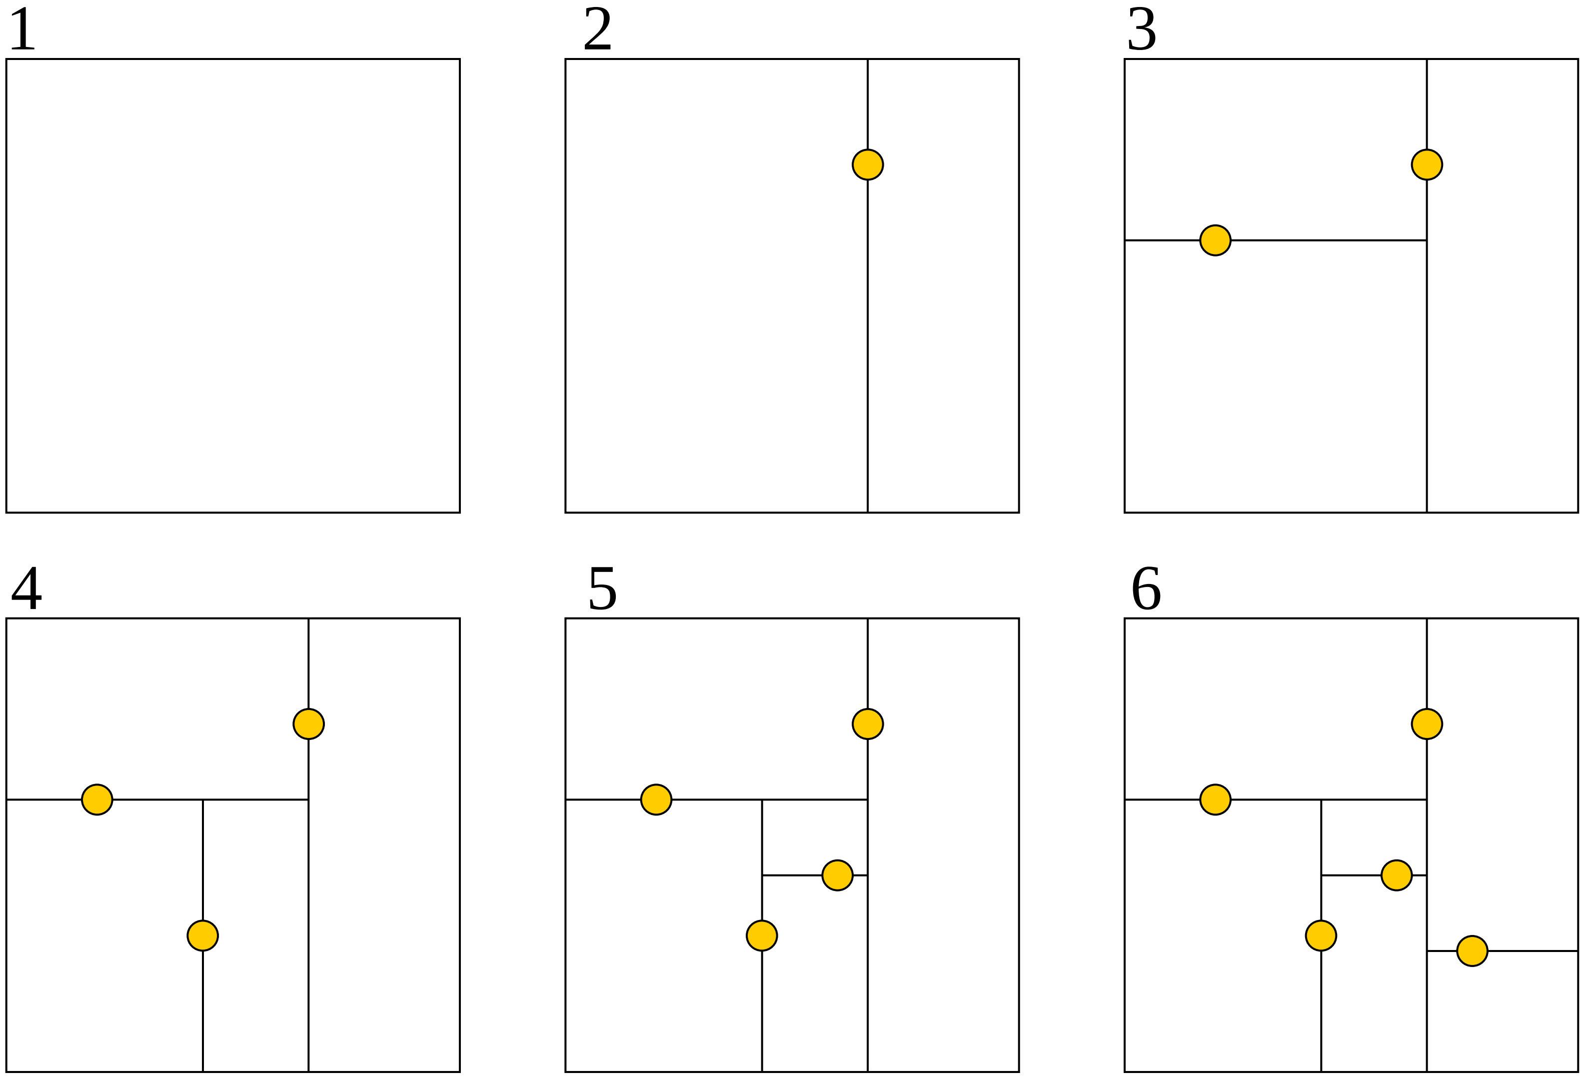

Pass 1: Photon Tracing

- Compact, point-based photons to propagate flux through the scene
- Photons are traced from the light sources and propagated through the scene just as rays are in ray tracing
  - i.e. they are reflected, transmitted, or absorbed
- To propagate photons use:
  - Russian roulette
  - Standard Monte Carlo sampling techniques
- For non-specular events, store them in the photon map

Pass 1: Photon Tracing - Caustics

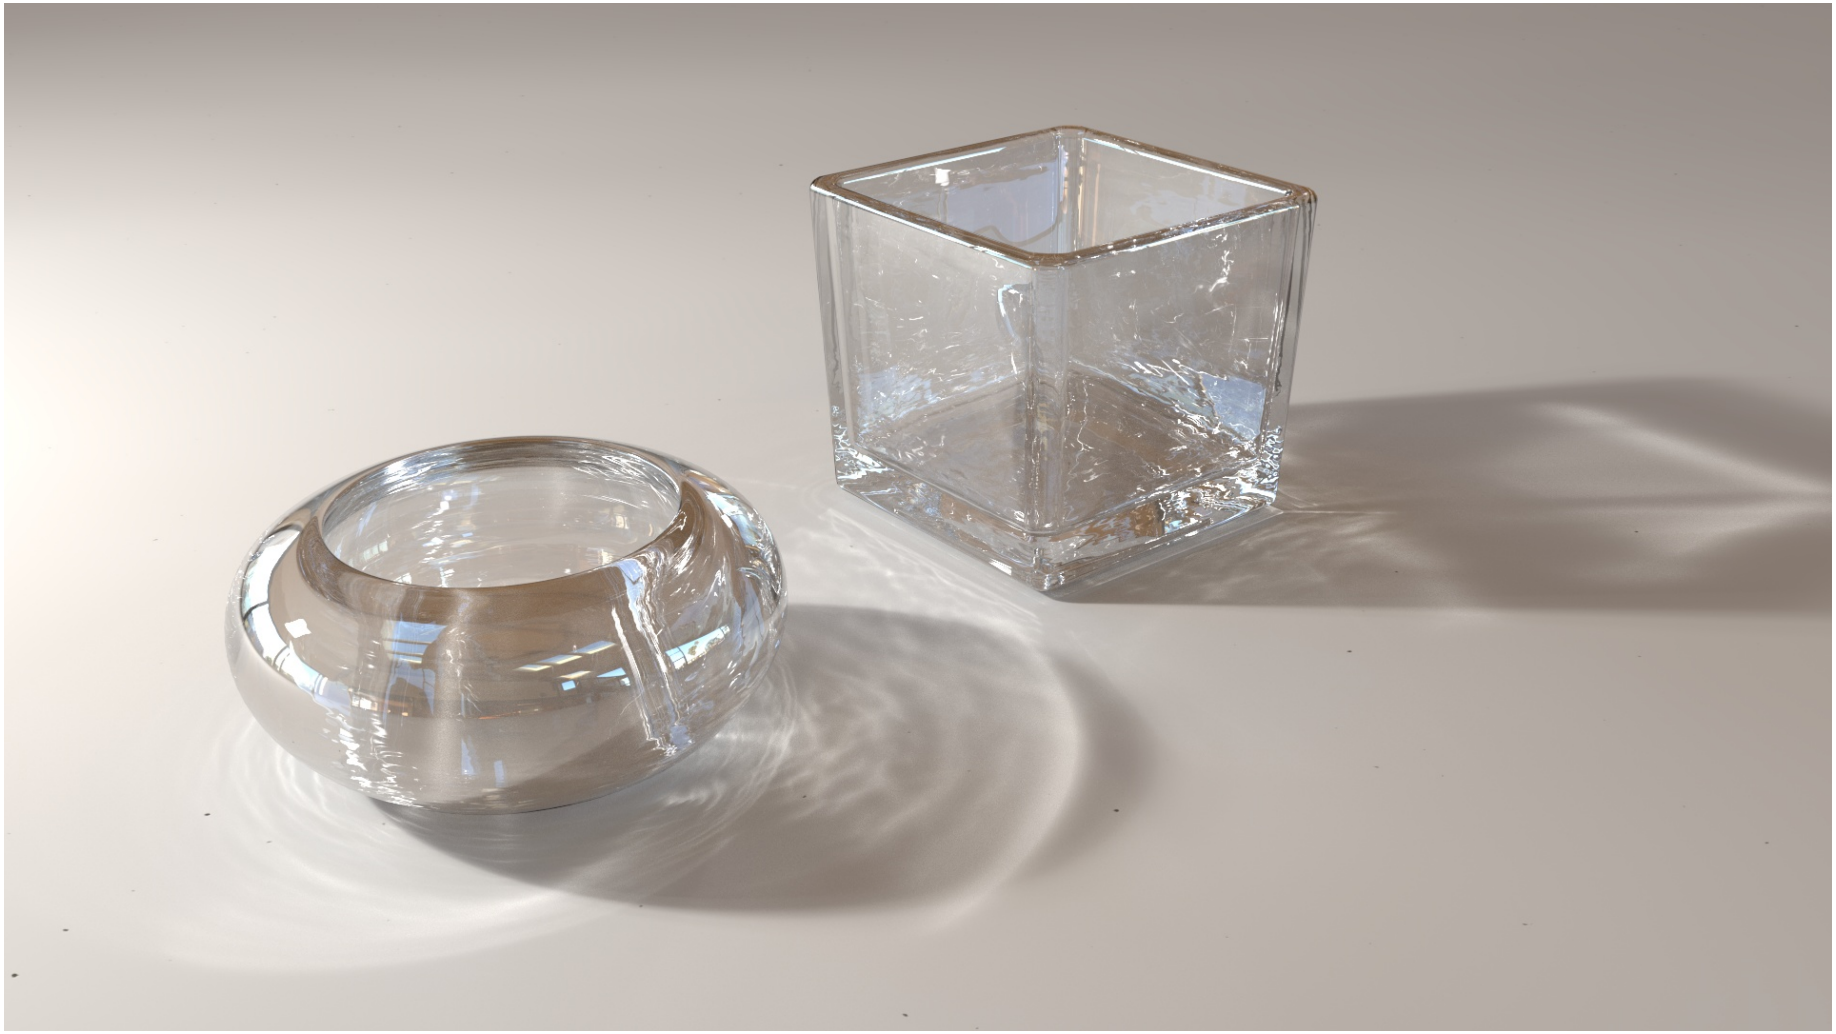

Pass 1: Photon Tracing - Caustics

- A caustic is formed when light is reflected or transmitted through one or more specular surface before reaching a diffuse surface
- Photon mapping can be efficient for computing caustics:
  - Photons for caustics need to be directed only at **specular surfaces→**
  - Are **fewer** than those for indirect illumination
  - Can be stored in a separate photon map: **caustics map**

Identity of a Photon (1)

- Two strategies:
- 1) Shoot generic photons and change identity along the path according to surfaces encountered
- 2) Shoot **global illumination and caustics** photons separately (two photon maps)
  - Still change label of photons on diffuse event for caustic rays
- The second is more efficient

Identity of a Photon (2)

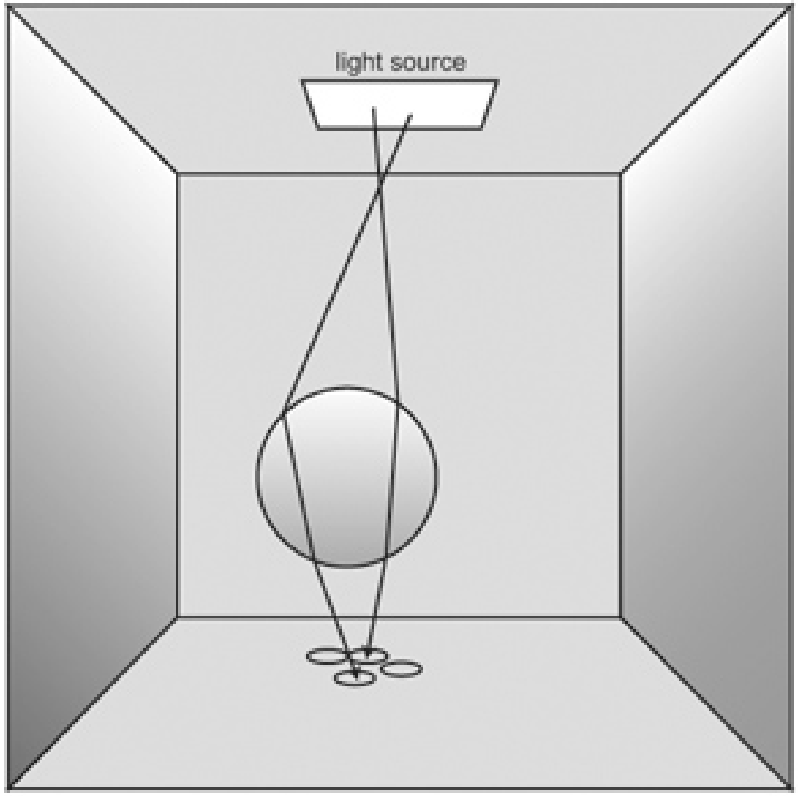

Caustic Map

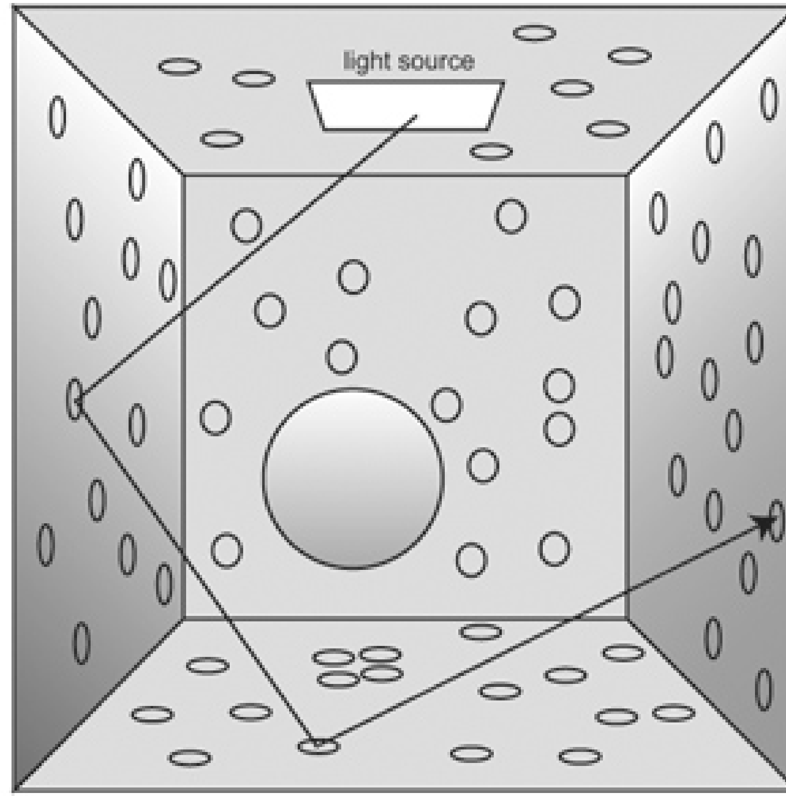

Global Photon Map

Identity of a Photon (3)

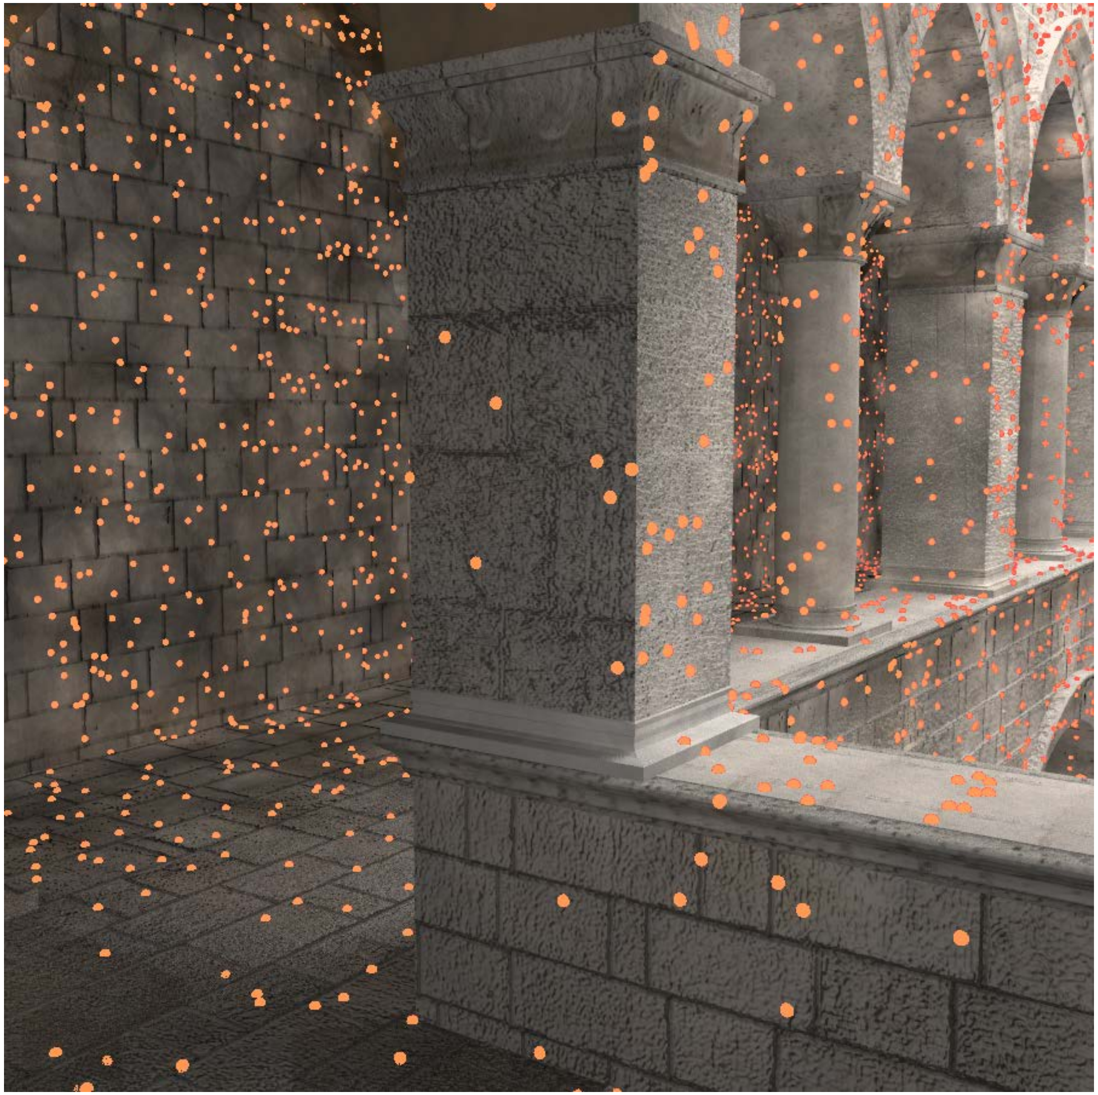

Global illumination photon map hit distribution

How to Cover all Lighting Effects

- Direct illumination (and shadows):
  - Monte Carlo source sampling from first camera intersection point
- Specular (mirror) reflection and refraction:
  - Plain ray tracing / stochastic path tracing
- Caustics and specular phenomena:
  - Tracing of caustics photons, directed at specular surfaces, storage in caustics map
- Global illumination and other effects:
  - Tracing of GI photons, storage in GI photon map

Pass 2: Rendering the Image

- Compute first image hits
- Calculate Direct illumination by Monte Carlo sampling
- Sample the caustics map to estimate outgoing radiance from caustics
- Sample the GI photon map to estimate outgoing radiance for indirect lighting
- Ray trace reflected and refracted rays (perform pass 2 for the newly discovered hits)

Photon Map Sampling

- Stored photons represent incident flux
- For an arbitrary point $\mathbf{x}$:
  - Locate N nearest stored points (k-d tree search)
  - Compute their flux contribution Φ_p according to distance from $\mathbf{x}$
  - Convert to radiance and find outgoing radiance
  - r = maximum distance to nearest particle

Photon Map Sampling (2)

$$L_r(\mathbf{x},\omega) = \int_{\Omega} f_r\left(\mathbf{x},\omega_p,\omega\right) L_i\left(\mathbf{x},\omega_p\right) \cos\theta_p \, d\sigma(\omega_p) =$$

$$L_r(\mathbf{x},\omega) = \int_{\Omega} f_r\left(\mathbf{x},\omega_p,\omega\right) \frac{d^2\Phi_i(\mathbf{x},\omega_p)}{\cos\theta_p \, d\sigma(\omega_p) dA(\mathbf{x})} \cos\theta_p \, d\sigma(\omega_p) =$$

$$L_r(\mathbf{x},\omega) = \int_{\Omega} f_r\left(\mathbf{x},\omega_p,\omega\right) \frac{d\Phi_i(\mathbf{x},\omega_p)}{dA(\mathbf{x})}.$$

$$L_r(\mathbf{x},\omega) = \sum_{p=1}^{n} f_r\left(\mathbf{x},\omega_p,\omega\right) L_i\left(\mathbf{x},\omega_p\right) \cos\theta_p \, \mathrm{d}\sigma(\omega_p) =$$

Photon Map Sampling (3)

$$L_r(\mathbf{x},\omega) \approx \sum_{p=1}^{n} f_r(\mathbf{x},\omega_p,\omega)\frac{\Delta\Phi_i(\mathbf{x},\omega_p)}{\Delta A(\mathbf{x})}$$

Assuming the surface to be locally flat around $\mathbf{x}$, we can approximate the $\Delta A(\mathbf{x})$ by the circle enclosing the n (nearest) photons:

$$L_r(\mathbf{x},\omega) \approx \frac{1}{\pi r^2}\sum_{p=1}^{n} f_r(\mathbf{x},\omega_p,\omega)\Delta\Phi_i(\mathbf{x},\omega_p)$$

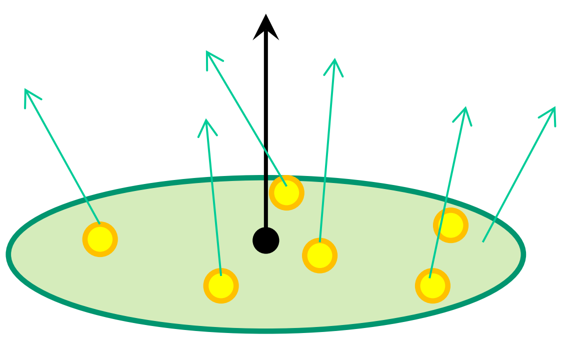

Photon Map Sampling – Examples (1)

- Caustics sampling:

Photon Map Sampling – Examples (2)

- Indirect diffuse illumination:

direct

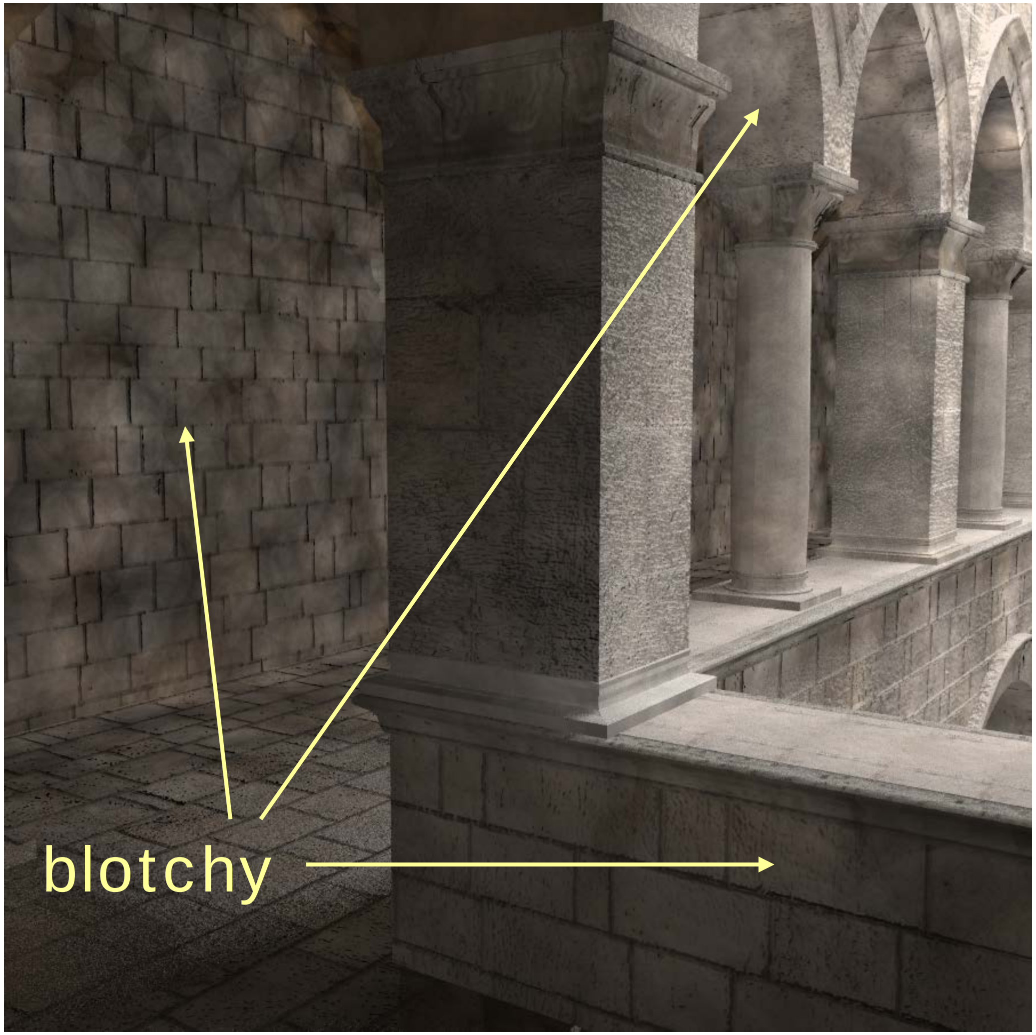

indirect

GI Gathering (1)

- To avoid visible artifacts caused by sparse GI particle distribution:
- Replace GI photon map sampling at the direct hit point by Monte Carlo sampling over the hemisphere of directions
- Gather GI from photon map at the next intersection step
- Essentially, shorten the GI photon tracing path by 1 bounce and
- Introduce stochastic sampling at final camera path node (+1 segment)

Caustics and GI Gathering Comparison

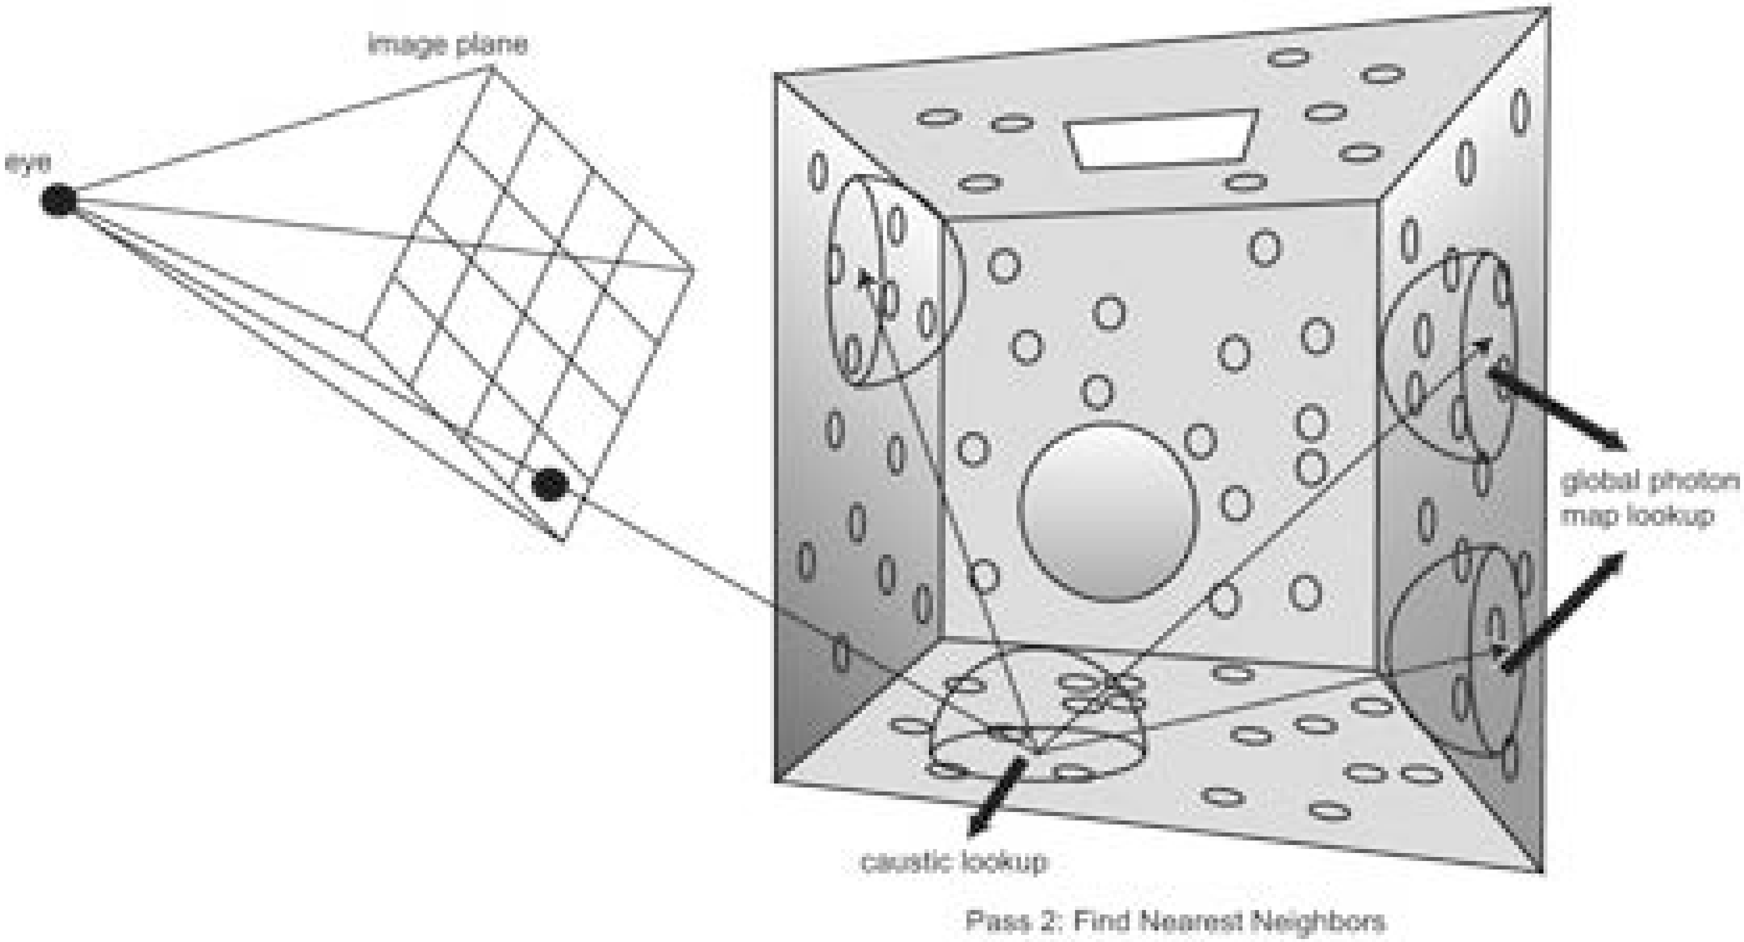

Gathering Example

No gathering

Gathering

Photon Mapping Example