Shadows

Georgios Papaioannou - 2017

FUNDAMENTAL CONCEPTS

Shadows and Perception (1)

- Wherever there is light, there are shadows
- Presence of shadows:
  - Not only for aesthetic purposes
  - Provides clues for the shape of the geometry in the image
    - Helps place the objects in the environment. Gives clues about relative distances
    - Enhances depth perception: In monocular vision the HVS relies on clues and recognizable configurations to discern the ordering and distances of objects
    - Indicates the direction of incident light or light sources
- Enhances the visual detail of the displayed surfaces by enhancing local contrast

Shadows and Perception (2)

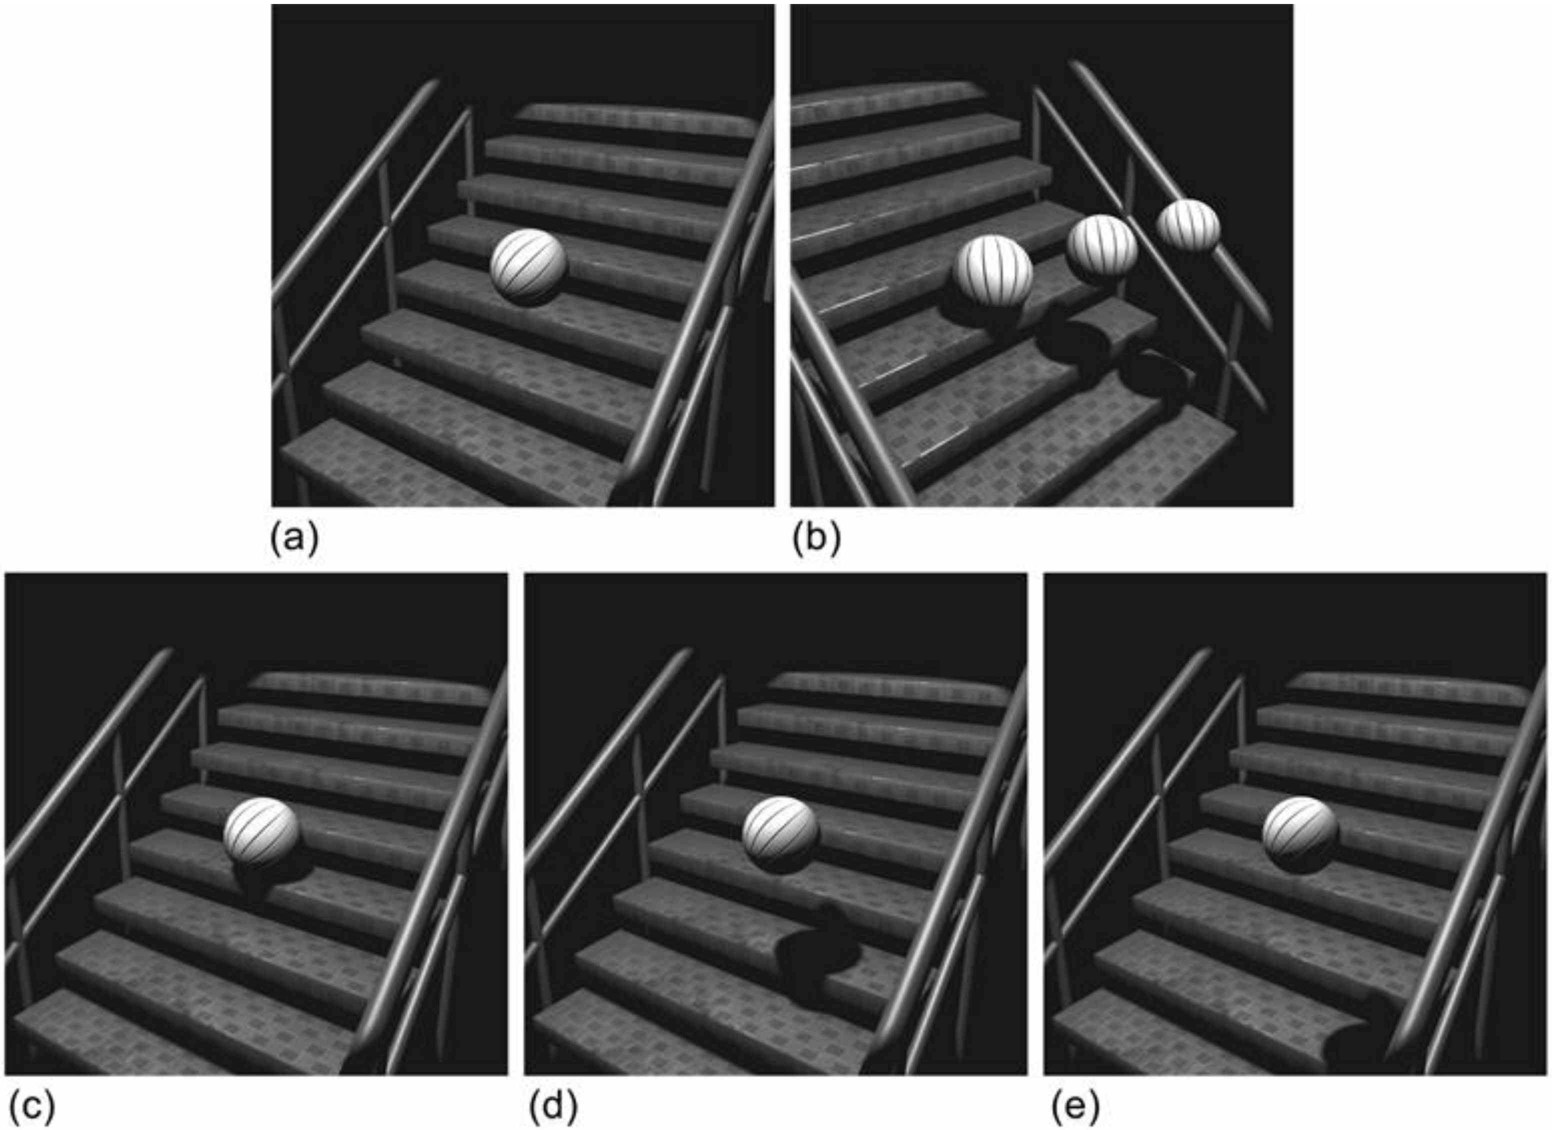

Shadows and Perception (3)

- (a) No shadow: We cannot possibly know the relative position or size of the ball w.r.t. the steps
- (b) Possible position/ball size configurations that lead to the same image (a)
- (c,d,e) The resulting images of the configurations in (b) when shadows are enabled

Shadows and Visual Detail

(no shadows)

Coarse, uninteresting surfaces

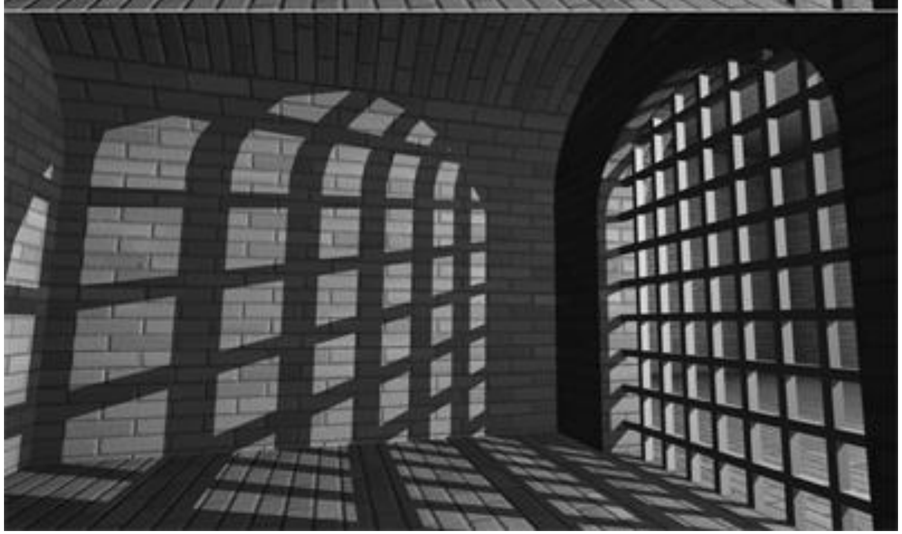

(with shadows)

Same geometry, higher visual detail

How are Shadows Generated?

- Partial or full obstruction of a source's light by geometry
- Indirect illumination reaching a surface is in general of lower luminance compared to the direct, unshadowed light →
- Illuminance of points in shadows is significantly lower than that of the lit points

Shadow Types

- The size and type of shadows depend on the size and distance of the light emitting surfaces:
  - Infinitely distant light (directional) sources cause parallel shafts o shadows
  - Non-directional light sources cause radially projected shadow profiles

Umbra and Penumbra

- Umbra is part of the shadow due to complete light obstruction
- Penumbra is the shadow part where partial occlusion occurs and creates a soft transition to the lit surface (soft shadows)
- A punctual (point) light source creates hard shadows with no penumbra
- A light source with a non-negligible size and comparable distance to the occluding geometry causes shadows with penumbrae (soft shadows)
  - Larger emitters and smaller distances to occluders → larger penumbrae

Shadow Examples

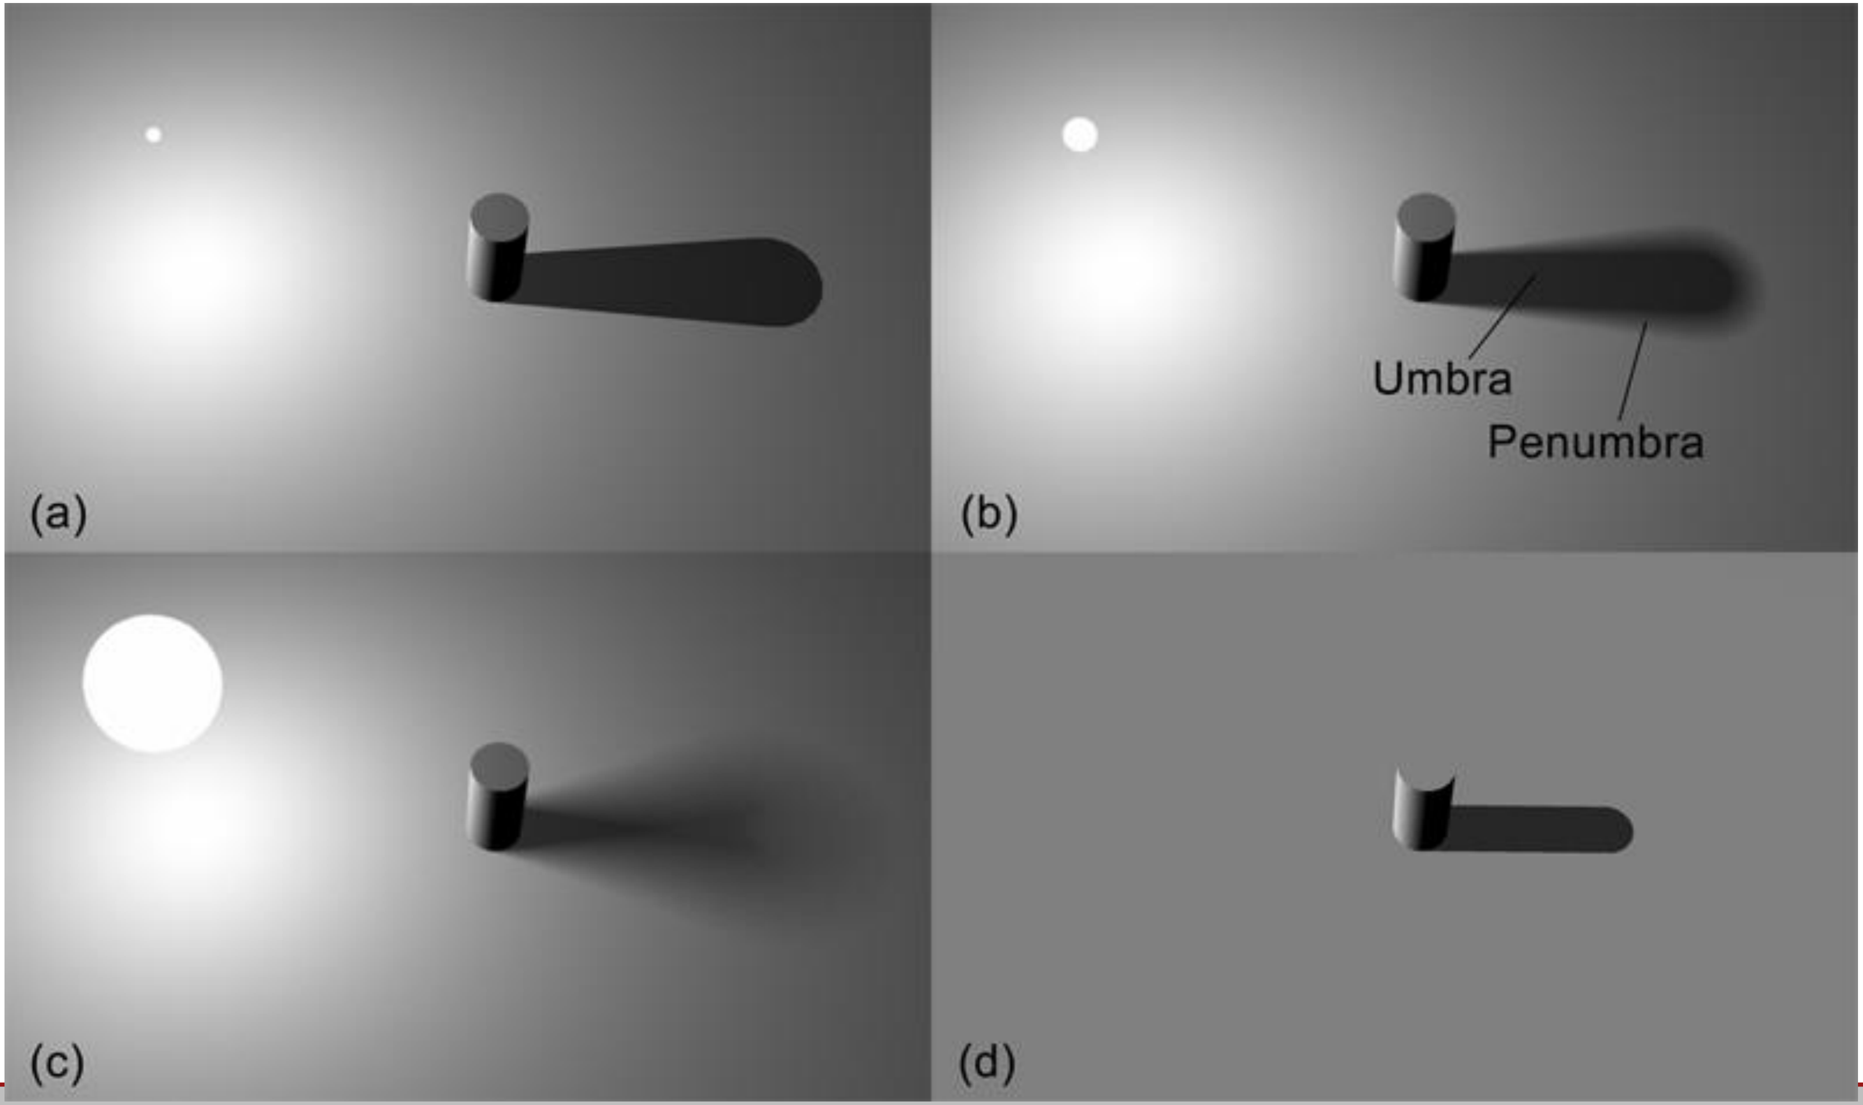

Shadow Volume

- The interception of the emitted light by an obstacle creates a volume frustum spanning the silhouette of the geometry and extending to infinity
- Any point within this volume is in shadow

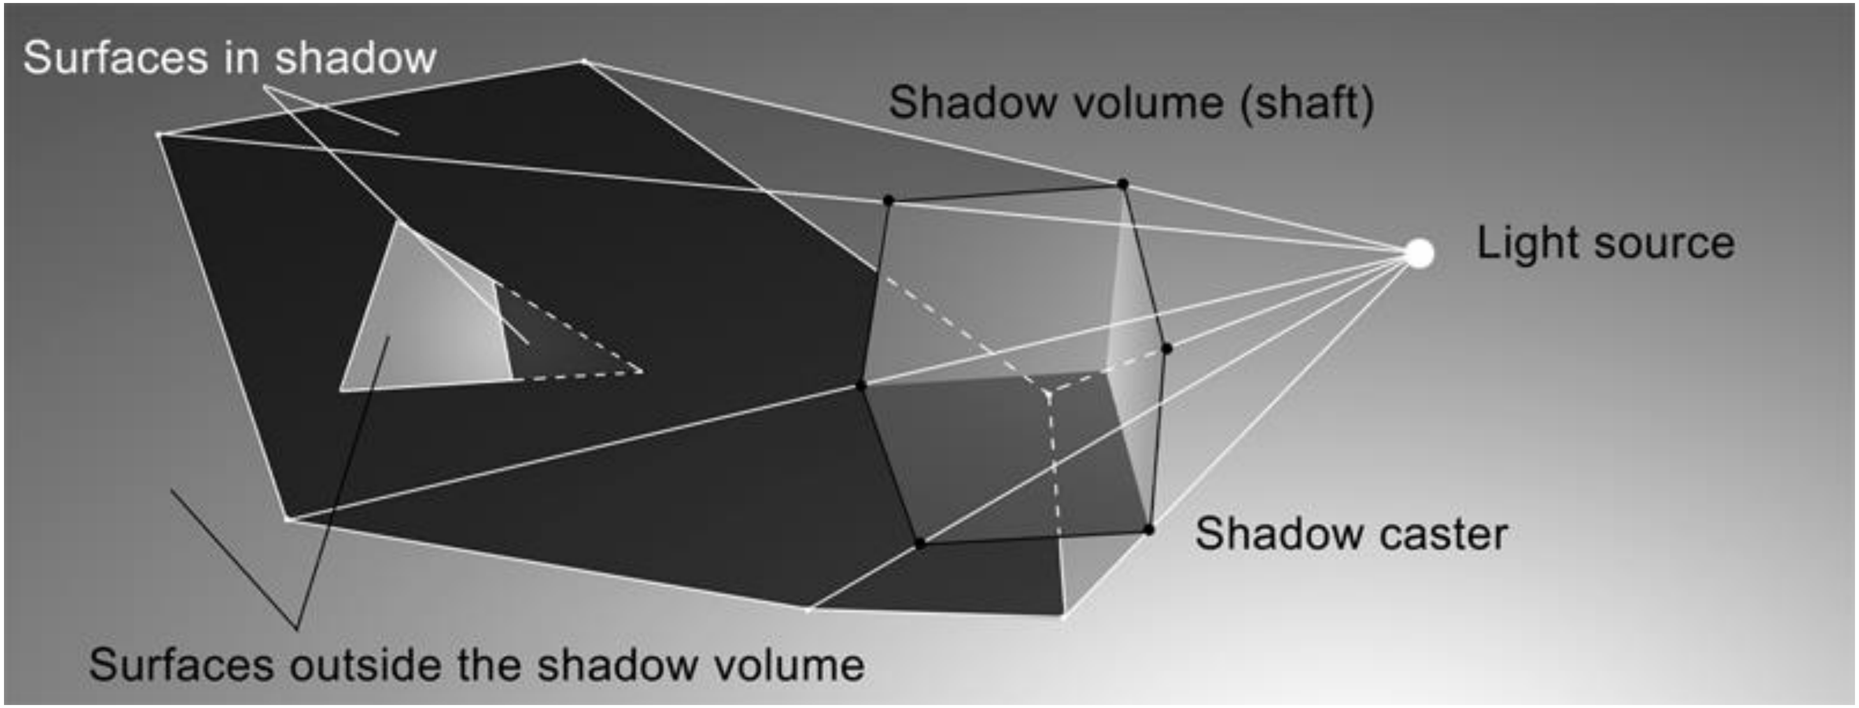

SHADOW GENERATION: SHADOW MAPS

Shadow Maps

- Basic principle:
  - The occlusion of light on a surface due to a given (point) light source is a similar problem to the visibility determination from the user's view point
  - A point is lit if the point is the closest one to the light source in this direction, i.e. if it is "visible" from the light source
- We can use the depth buffer mechanism to perform HSE and determine the nearest visible points from the light source's view point
- We call the depth buffer generated from the light source view point a shadow map

Shadow Map - Setup

- A projection is set up from the light source's point of view (a) and the shadow map is captured (b)
- The scene is rendered normally form the camera view point and fragments are tested against the shadow map (c)

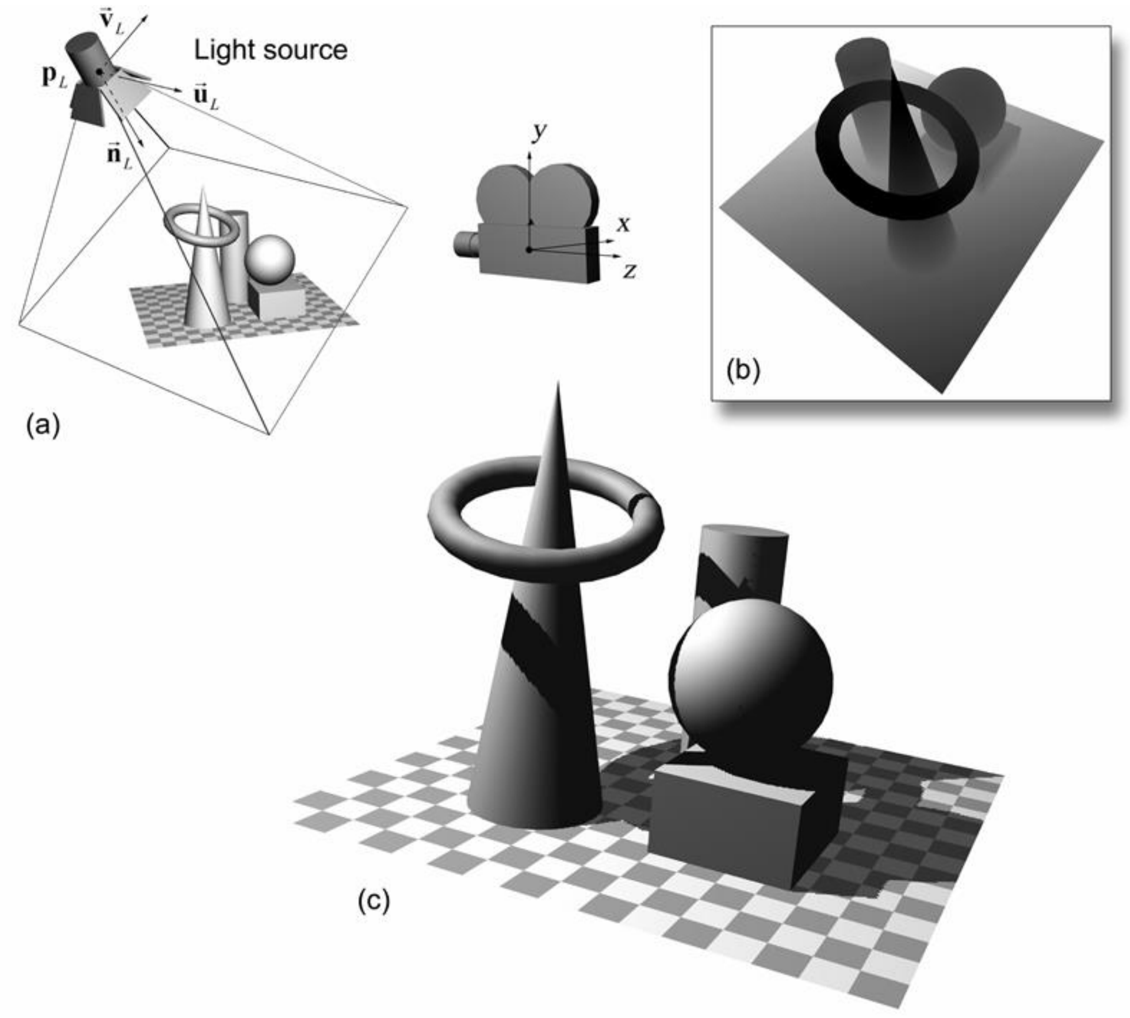

Transforming Fragments to S.M. Space

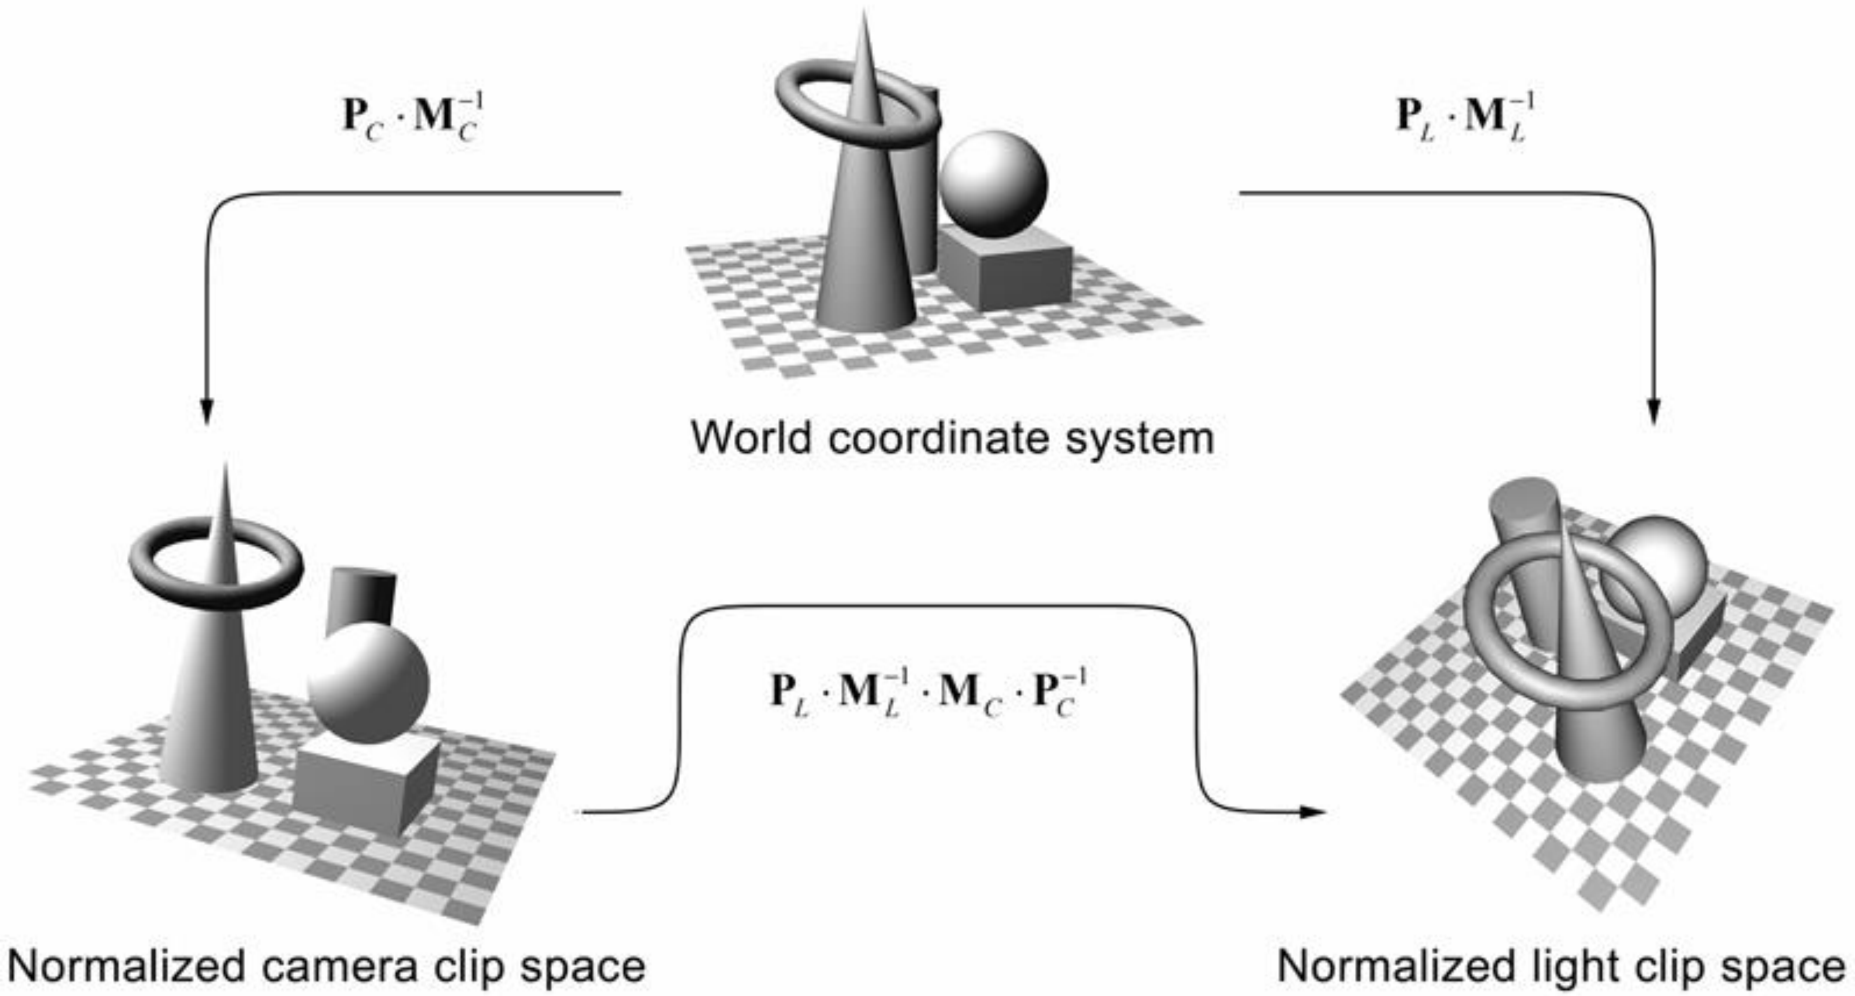

Shadow Calculations

- Render the scene from the light source view point $(\mathbf{p}_L, \vec{\mathbf{u}}_L, \vec{\mathbf{v}}_L, -\vec{\mathbf{n}}_L)$
  - Transform geometry by $\mathbf{P}_L \mathbf{M}_L^{-1}$
  - Record the depth (shadow) map Z_L
- Render the scene normally, from the camera view point
  - Transform each fragment from the camera CSS to the light source's CSS:

$$\mathbf{p}'_{frag} = (x'_{frag}, y'_{frag}, z'_{frag}) = \mathbf{P}_L \cdot \mathbf{M}_L^{-1} \cdot \mathbf{M}_C \cdot \mathbf{P}_C^{-1} \cdot \mathbf{p}_{frag}$$

  - Compare the fragment's light space z'_{frag} value with the corresponding depth in the shadow map $Z_L(x'_{frag}, y'_{frag})$
  - If $z'_{frag} \leq Z_L(x'_{frag}, y'_{frag})$ the fragment is lit, otherwise it lies in shadow

Shadow Maps – Remarks (1)

- The shadow map needs to be updated only if:
  - The light source is moving
  - Geometry within the light's field of view changes
- The shadow map rendering time is significantly lower than the normal rendering time:
  - Only fragment depth is captured
  - No pixel shading occurs (pass through shader), no color attachment

Shadow Maps – Remarks (2)

- WYSIWYG: Whatever geometric entity can be rasterized or otherwise drawn in a depth map, can be used as an occluder:
  - E.g. foliage modelled as polygons with transparent textures

Advantages of Shadow Maps

- A simple and intuitive 2-pass algorithm
- Any renderable entity can generate shadow
- Easily combined with other effects, such as volumetric lighting
- Low complexity, takes advantage of GPU's early culling mechanisms
- Linear dependence on scene complexity
- Adjustable SM size → performance/quality trade off
- Can generate soft shadows (via extra samples)

Shadow Map Problems (1)

- Only works for conical/directional light sources
  - For omnidirectional lights, we need a cube map configuration of shadow maps

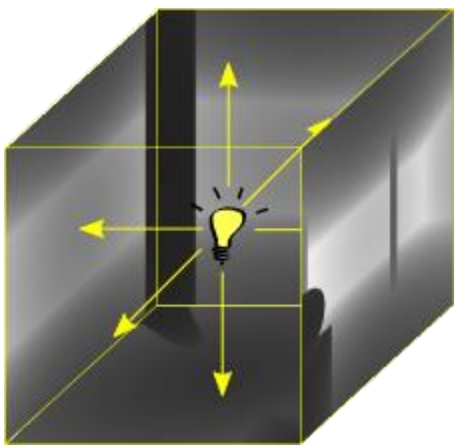

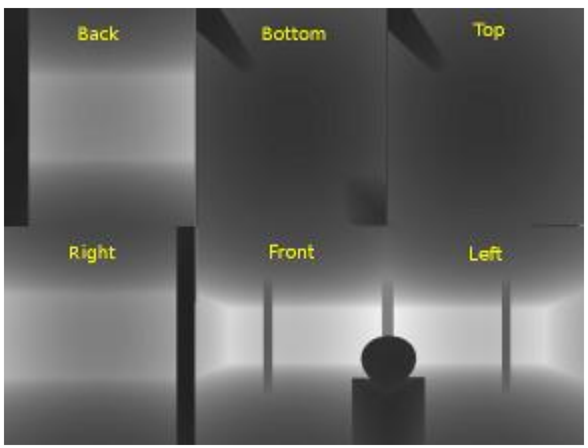

http://devmaster.net/p/3002/shader-effects-shadow-mapping

Shadow Map Problems (2)

- Accuracy depends on relative light-camera position and orientation
- Strong aliasing artifacts due to undersampling and arithmetic precision

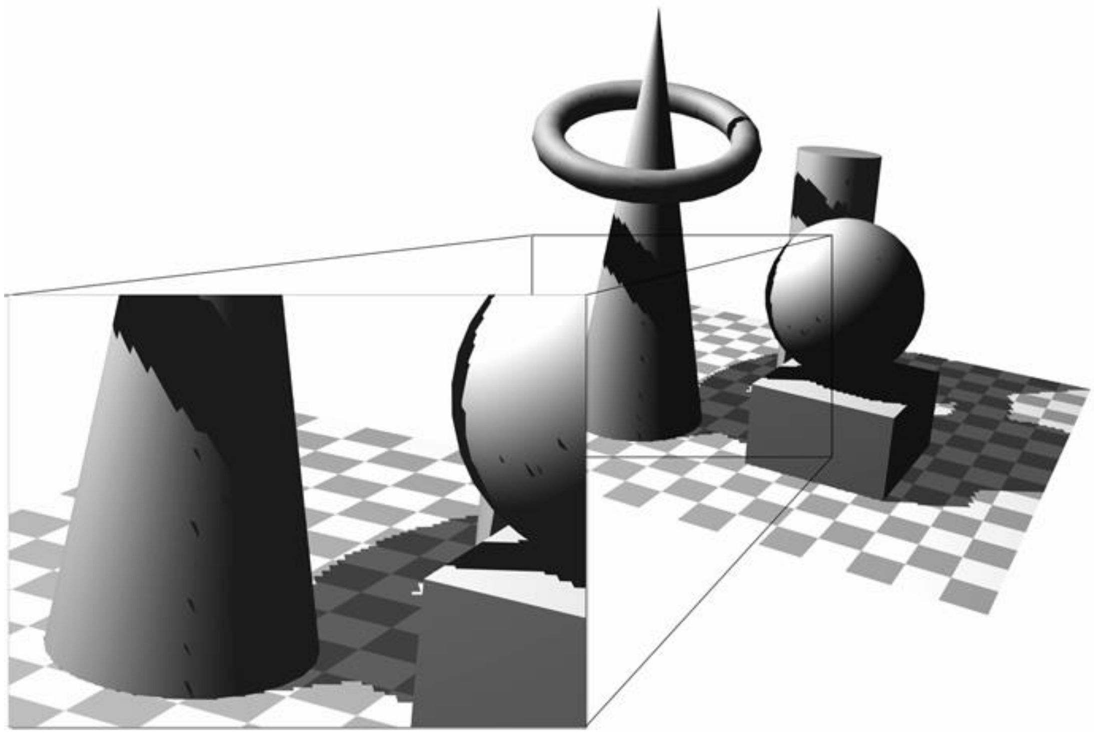

Typical Shadow Map Artifacts

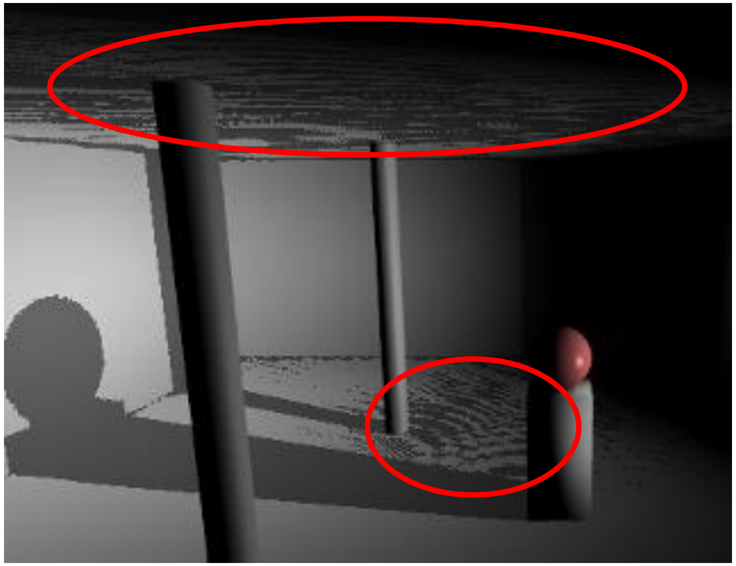

Shadow “acne”

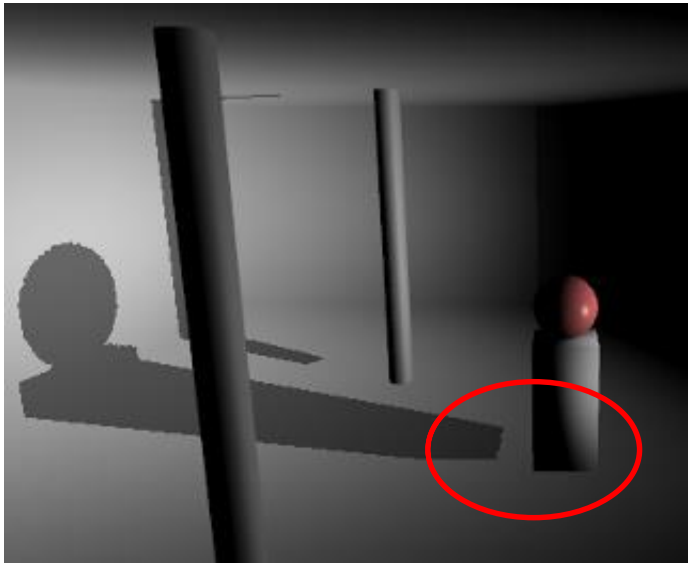

“Peter Panning”

Shadow Map Antialiasing

- Typical bilinear filtering on the shadow map does not work
- If we pre-filter (mipmap) the shadow maps:
  - We filter depths! → Erroneous depth comparisons and we do not get rid of artifacts
- We need to change the order of filtering and comparisons: post-filtering

Percentage Closer Filtering

- Draw samples from the shadow map in the neighborhood of the query shadow map coordinate
- Individually test each shadow map tap with the fragment z
- Average the shadow test results to get the fraction of occlusion

Percentage Closer Filtering

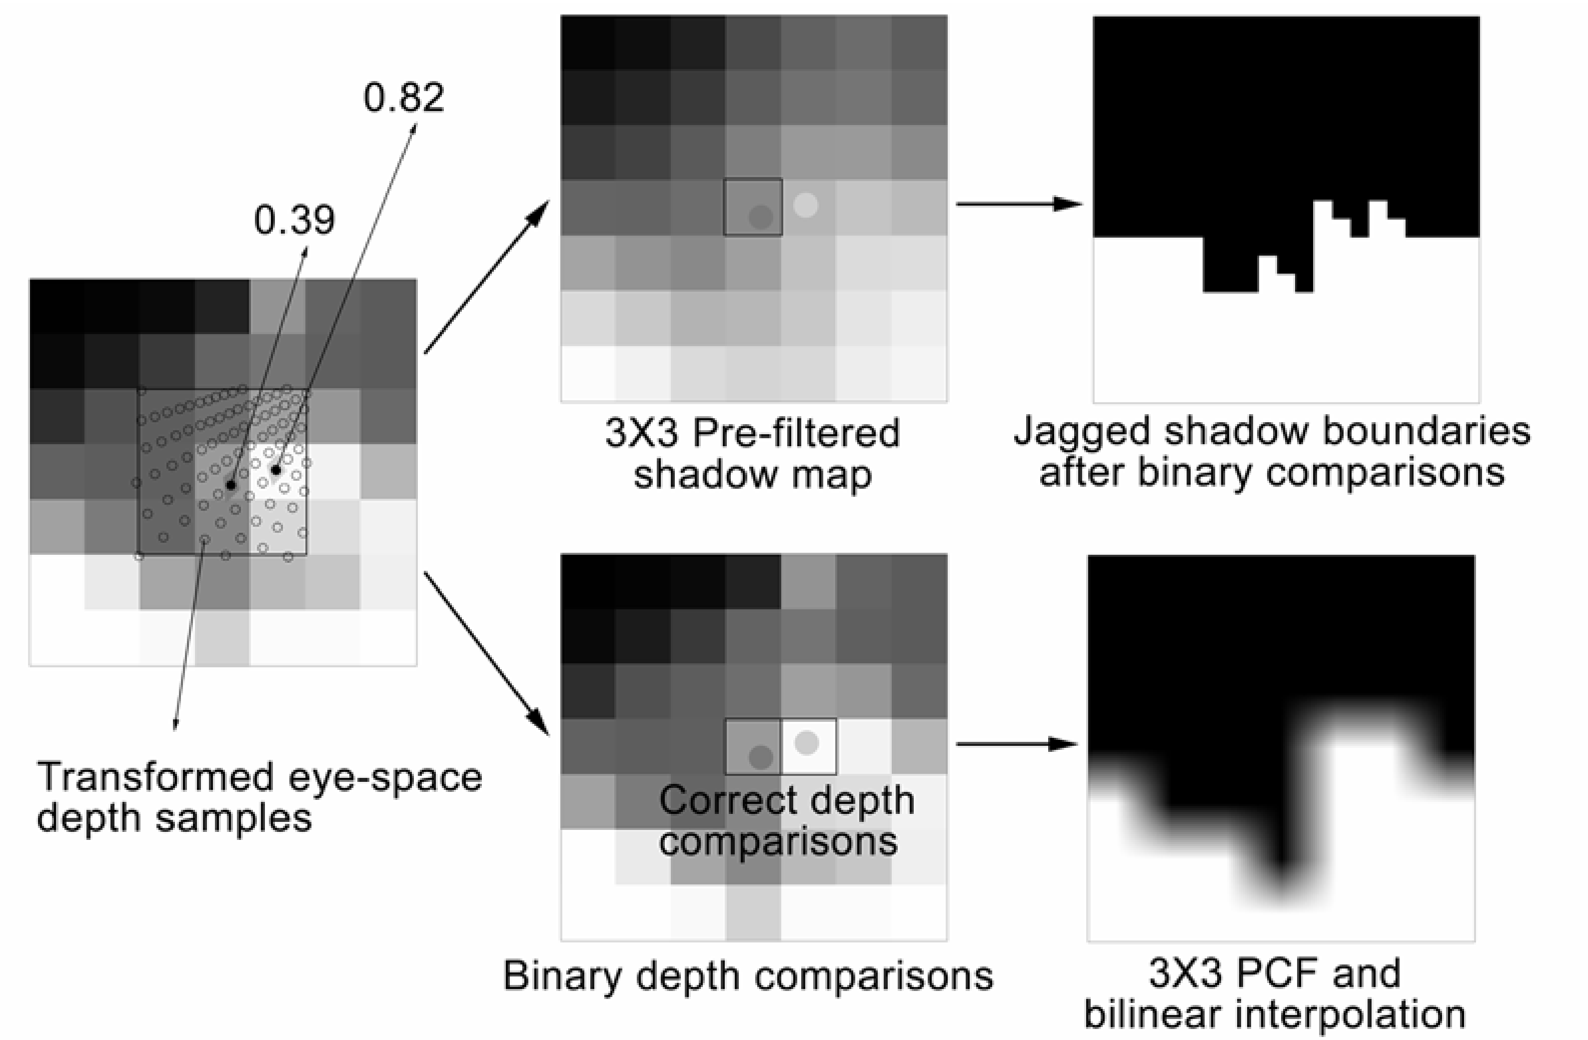

PCF Shadow Maps Example

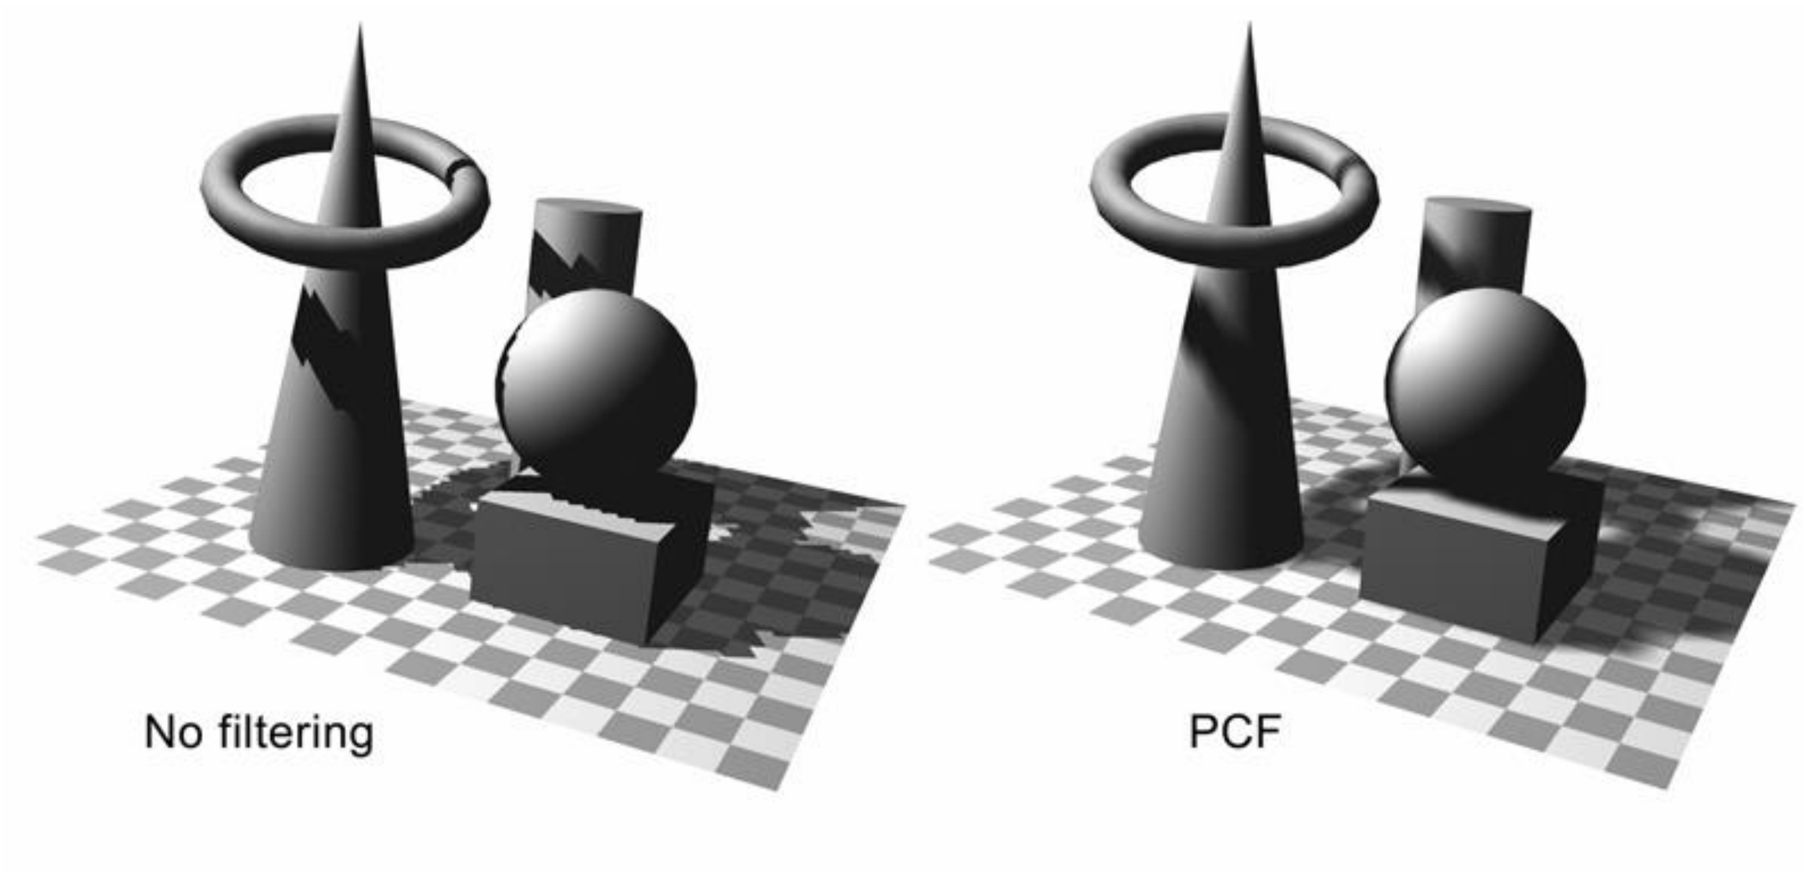

Contributors

- Georgios Papaioannou

References

[TP*06] T. Theoharis, G. Papaioannou, N. Platis, N. M. Patrikalakis, Graphics & Visualization: Principles and Algorithms, CRC Press