The Epipolar Occlusion Camera

Paul Rosen

Voicu Popescu

Department of Computer Science
Purdue University

Depth images

- Expressive
- Inexpensive to render
- Provide good LOD control

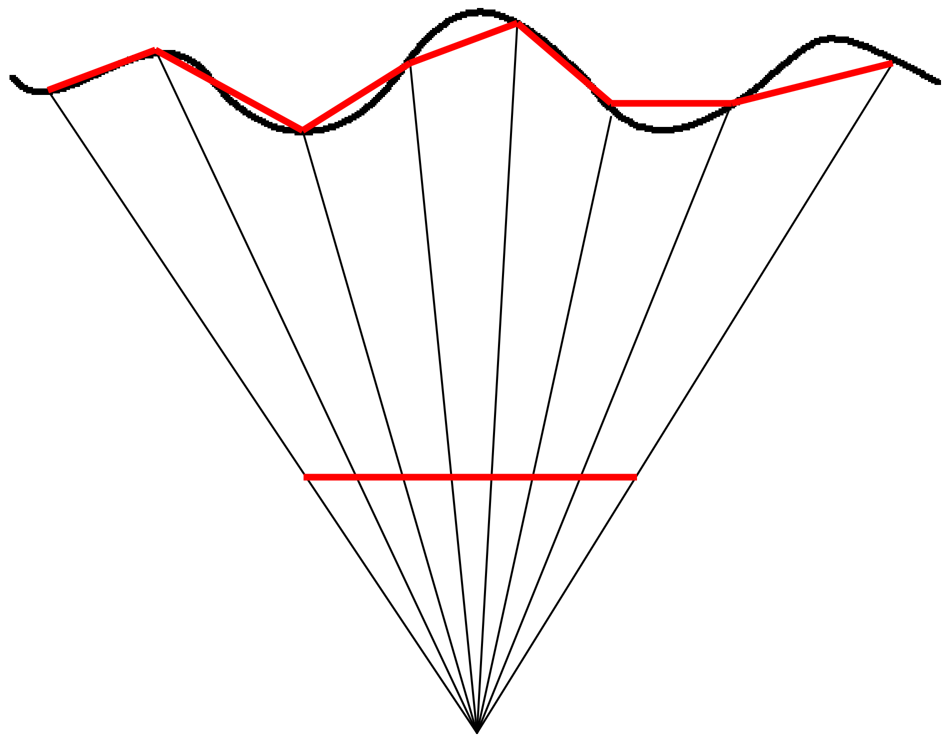

Depth image limitation

- Sample only what visible from reference viewpoint

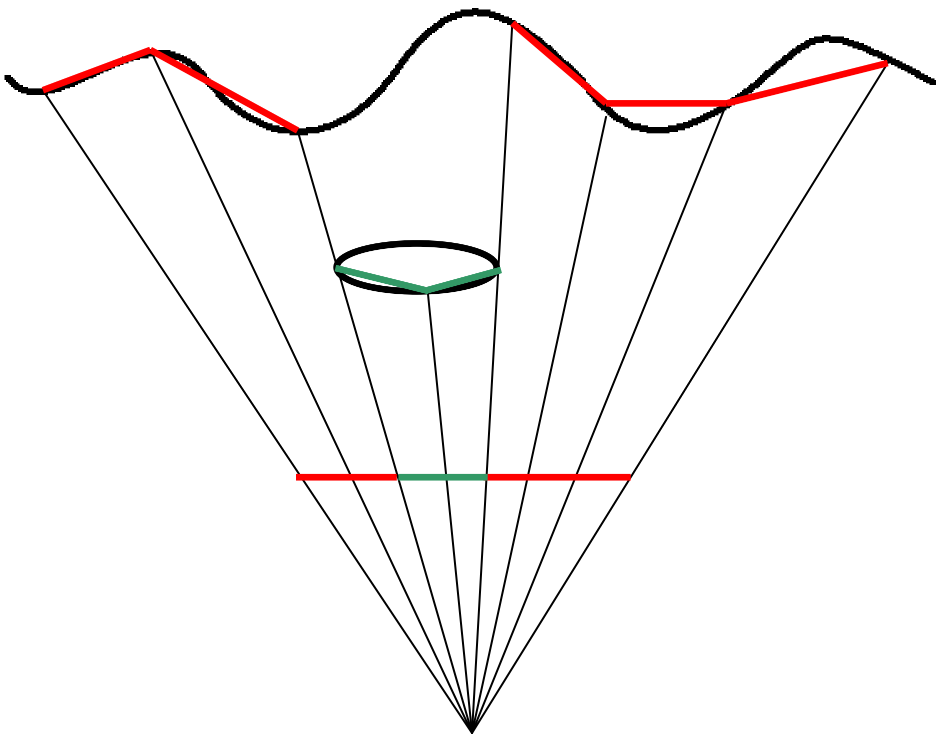

Depth image limitation

- Sample only what visible from reference viewpoint
- Missing samples cause gaps, or "disocclusion errors"

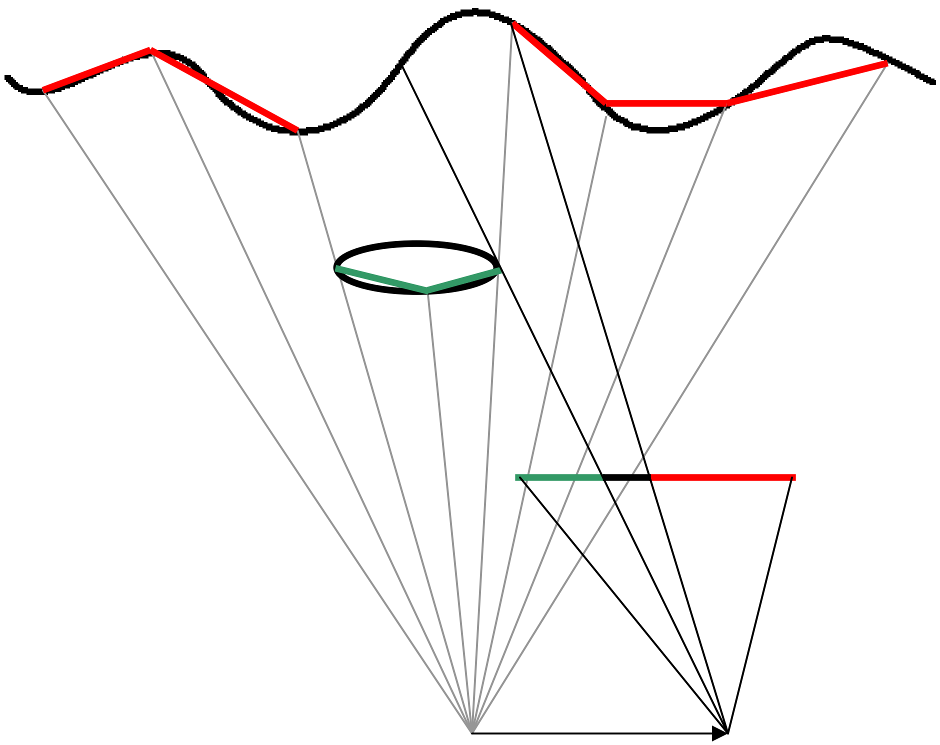

Disocclusion errors

- Even small viewpoint translations cause substantial artifacts

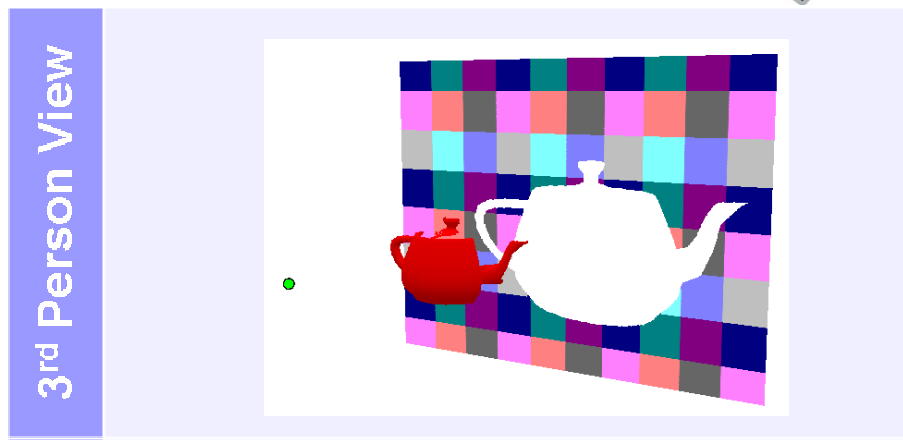

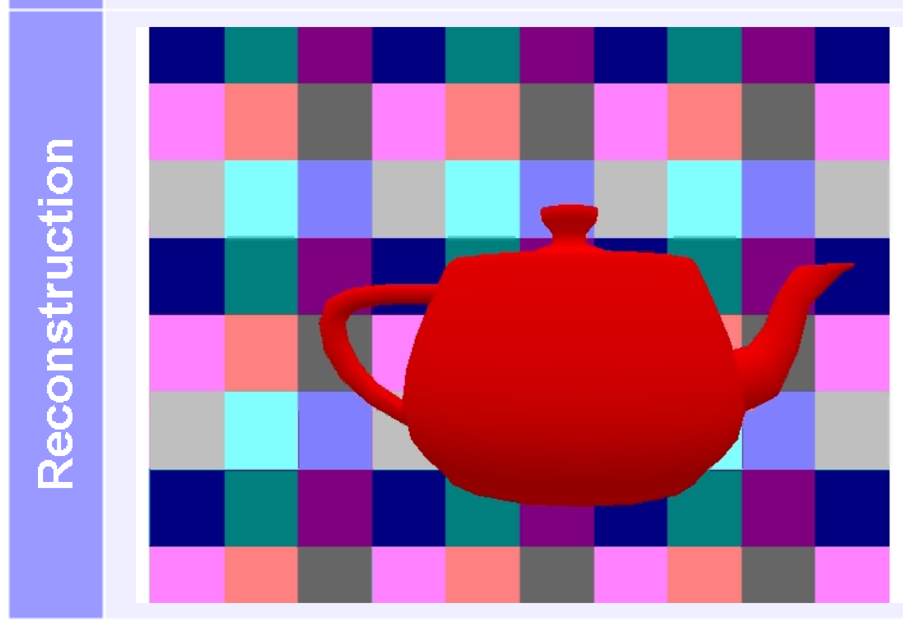

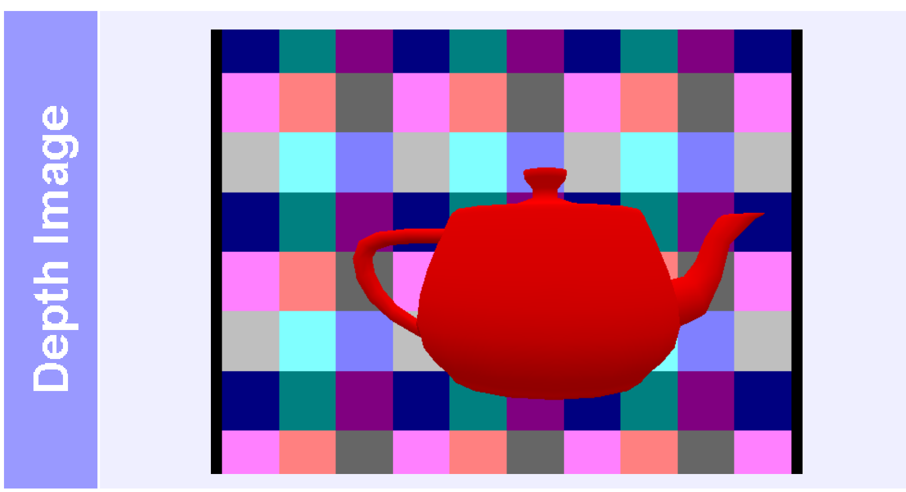

A palliative solution

- More depth images
  - Difficult to pick viewpoints
  - Many duplicate samples
  - Expensive

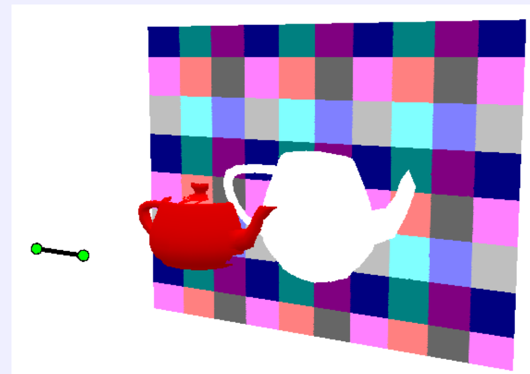

3rd Person View

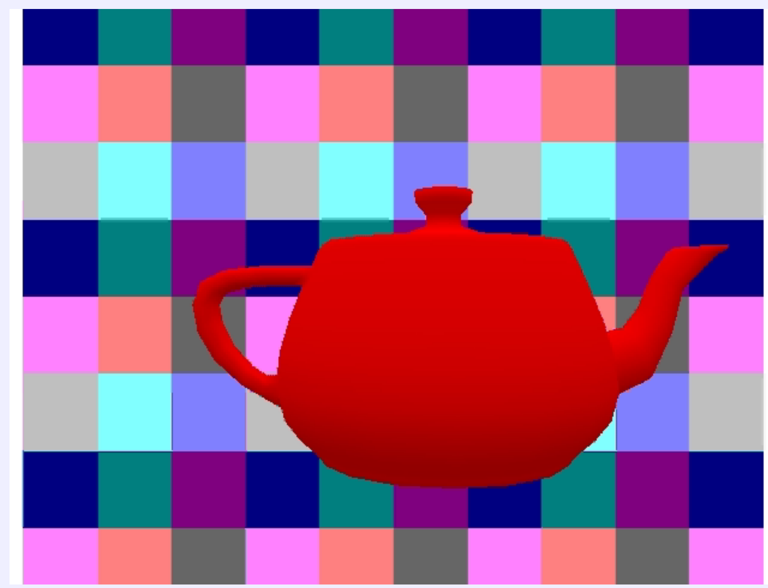

Reconstruction

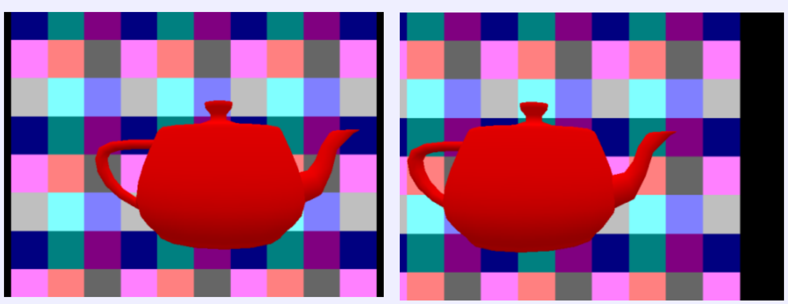

2 Depth Images

A solution: the EOC

- Render depth image with *epipolar occlusion camera*
  - EOC: non-pinhole that gathers all samples needed as the viewpoint translates

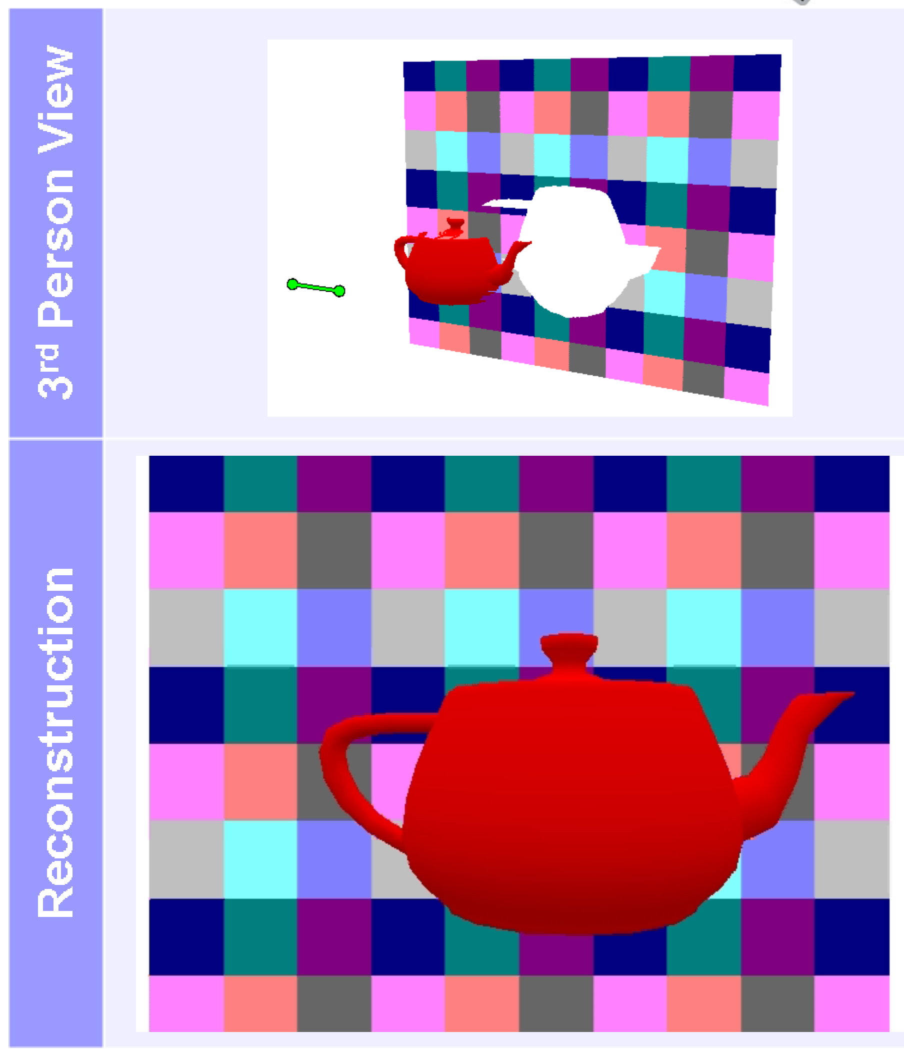

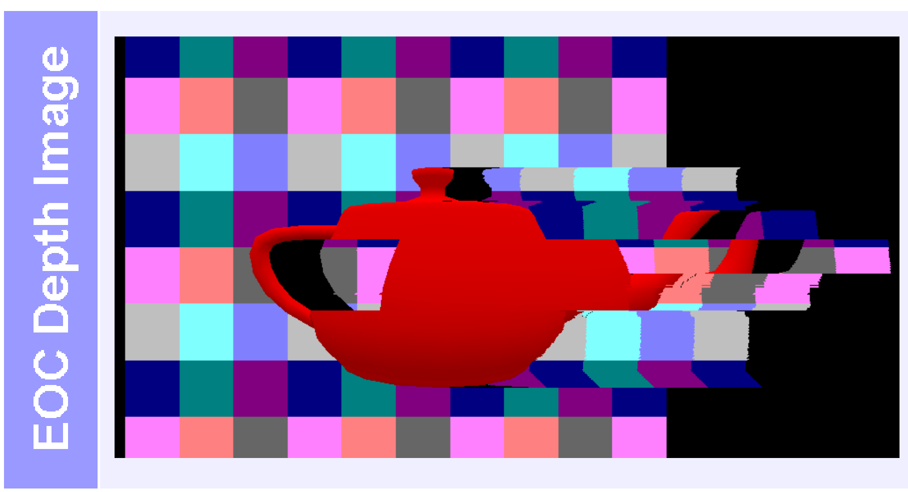

Outline

- Prior work
- EOC Construction
- EOC Rendering
- Results, discussion and future work

Occlusion cameras

- A family of cameras whose rays reach around occluders to gather samples that are barely hidden from the reference viewpoint thus needed to avoid disocclusion errors when the viewpoint translates

Prior work SPOC [Mei '05]

	Depth Images	3rd Person Views	
Pinhole	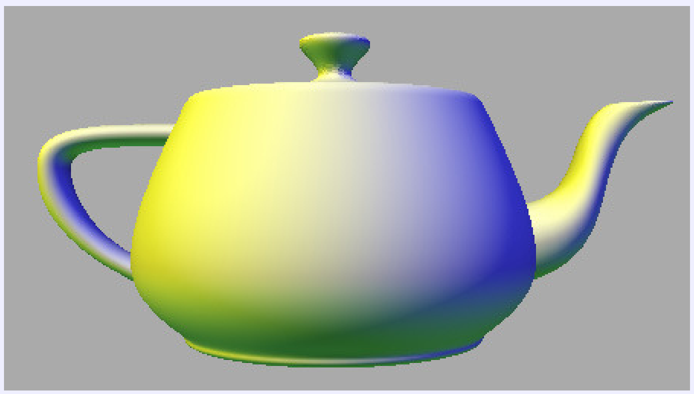		
SPOC	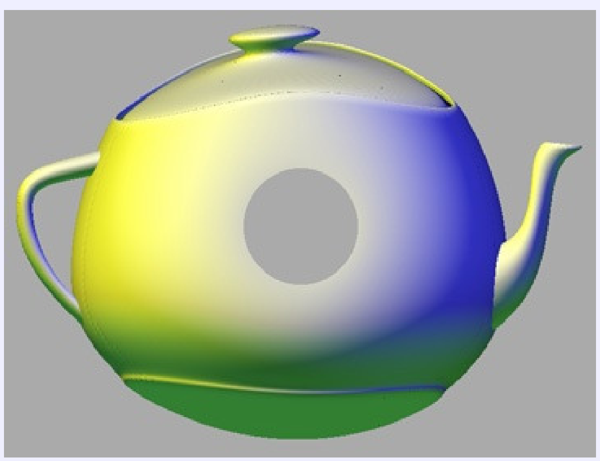		

Prior work DDOC [Popescu '06, Mo '07]

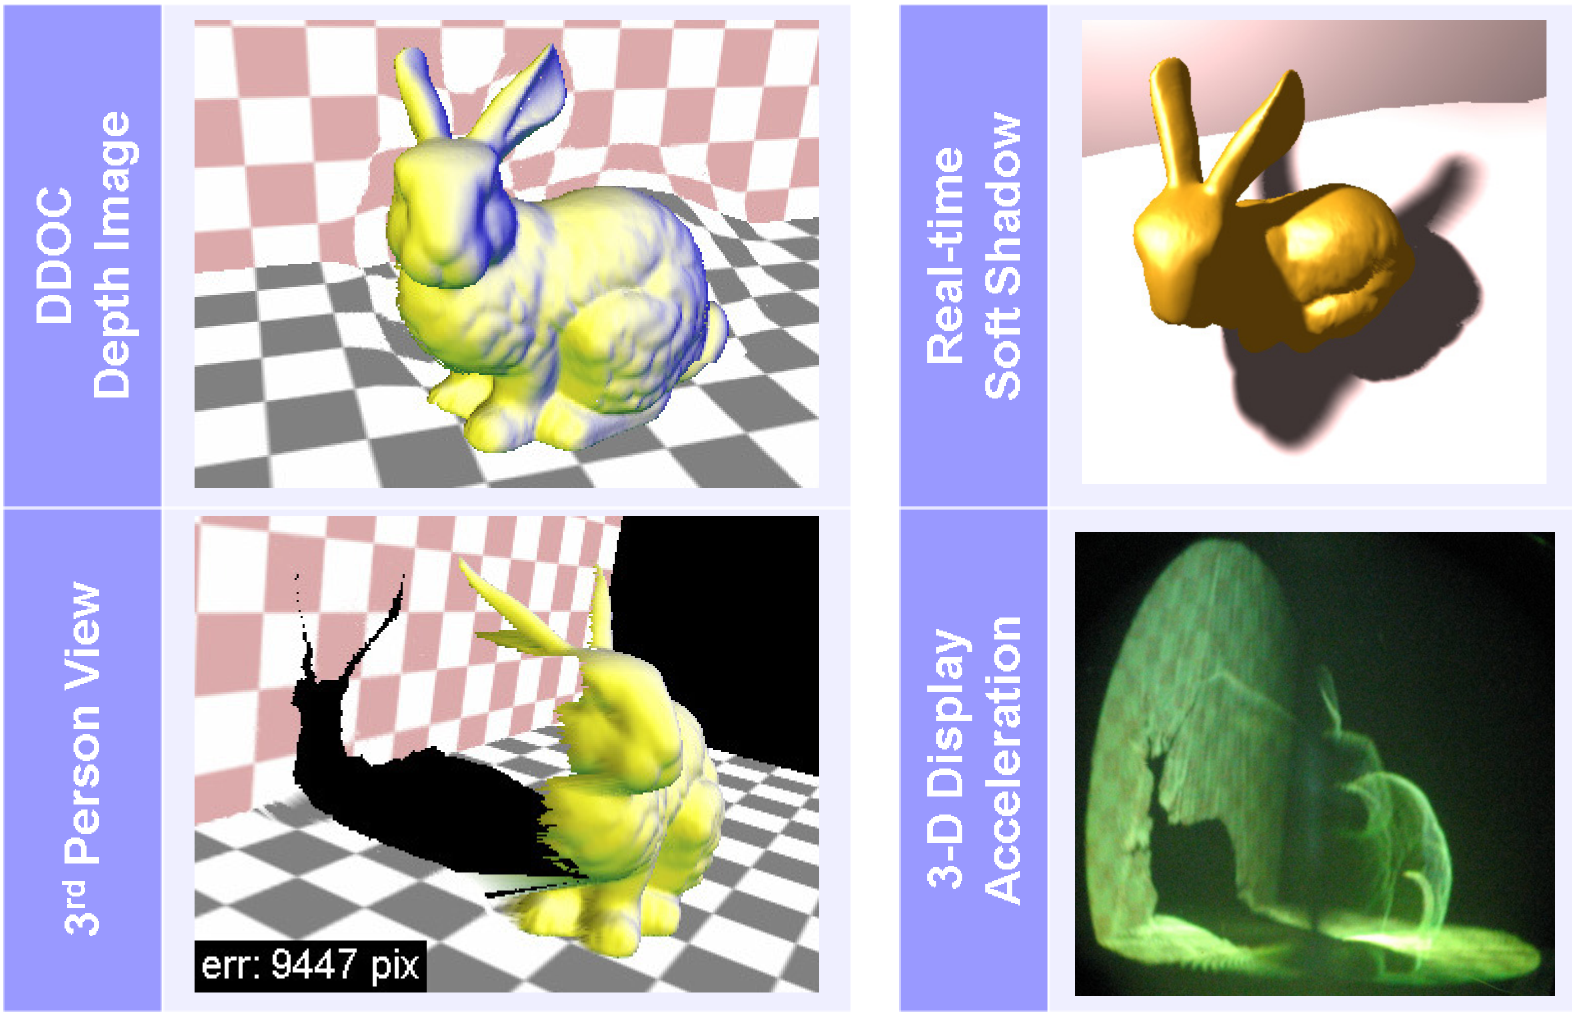

DDOC limitations

- Limited disocclusion capability for complex scenes
  - Depth discontinuities compete for same image region
- Set of novel viewpoints for which DDOC image has sufficient samples is defined empirically

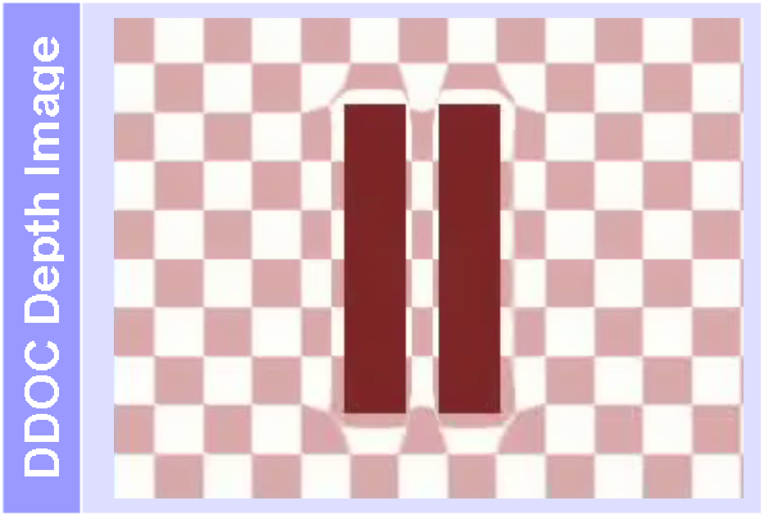

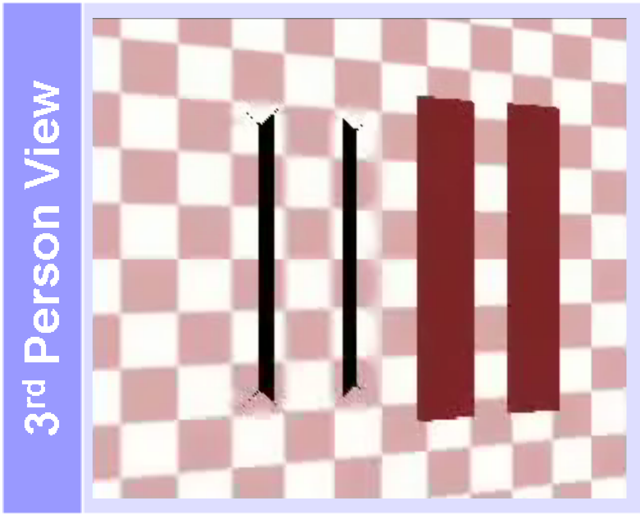

Epipolar occlusion camera

- View*segment* as opposed to view*point*
- Makes additional image space as needed to reveal occluded samples

EOC Depth Image

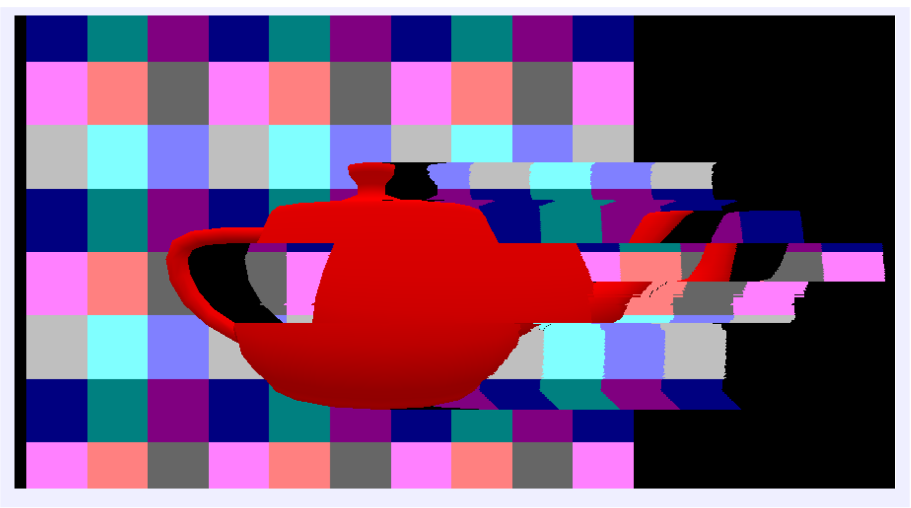

EOC Reconstruction

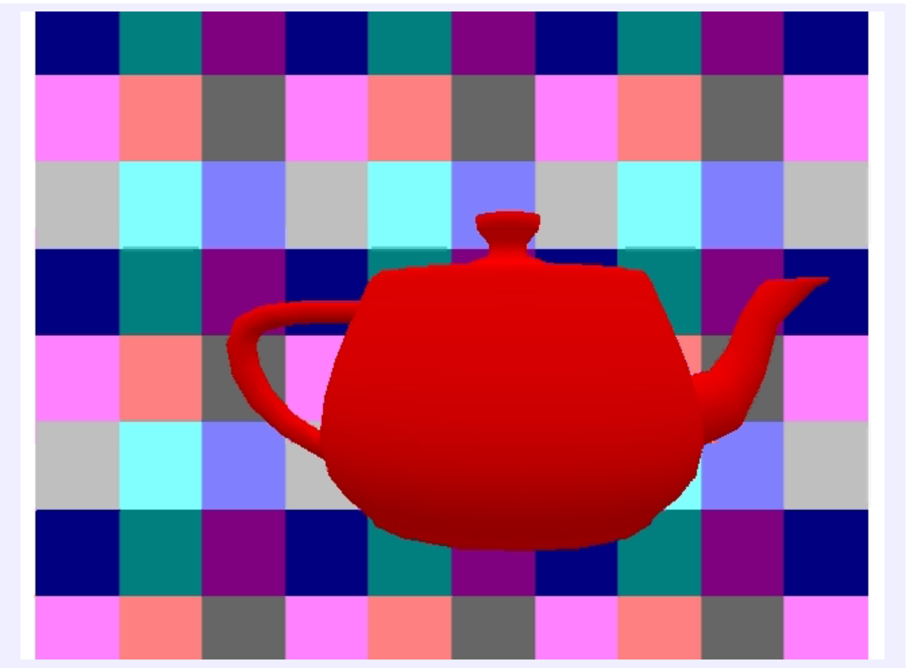

Construction

- Given
  - A 3-D scene
  - Two adjacent viewpoints *L* and *R*
- Construct camera model (i.e. set of rays) that is
  - *Comprehensive*—samples surfaces visible from viewpoints on segment *LR*
  - *Non-redundant*—rays do not intersect so a 3-D point has a unique projection
  - *Efficient*—the scene is rendered interactively

Epipolar constraints

- As viewpoint translates between L and R, samples "move" on epipolar lines
- Epipolar lines allow partitioning the disocclusion events in disjoint 1-D events
- The EOC model is constructed one epipolar line at the time

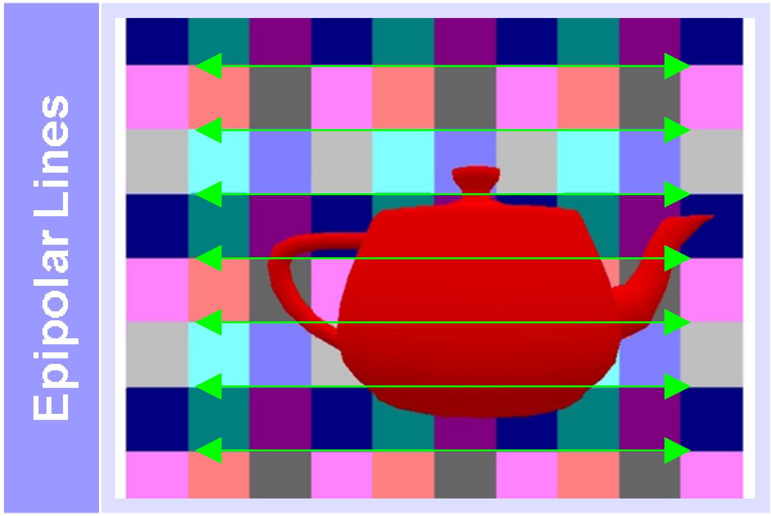

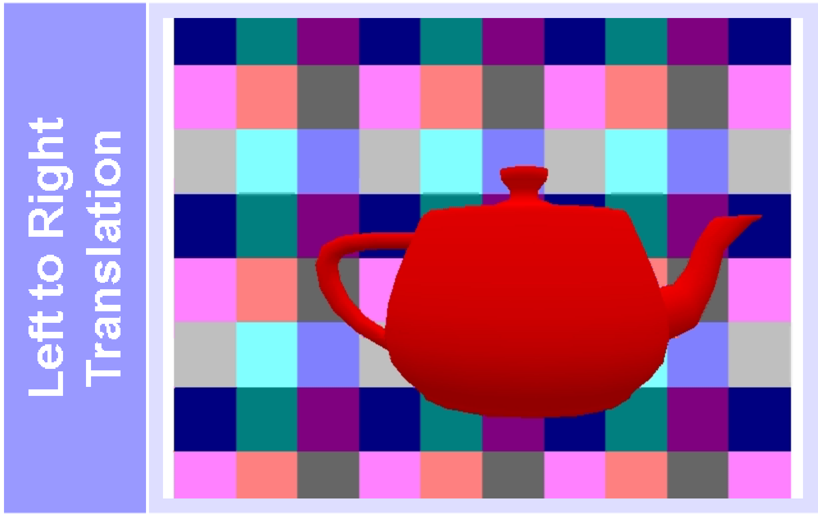

Construction

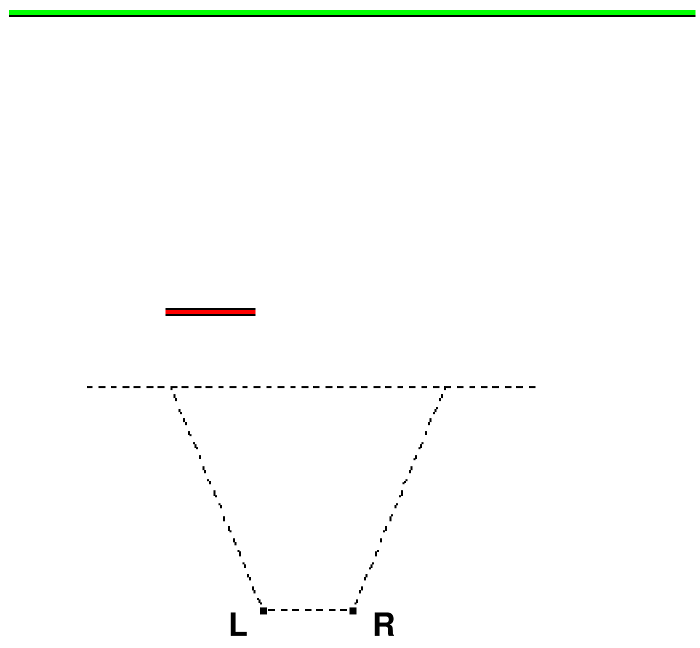

Construction

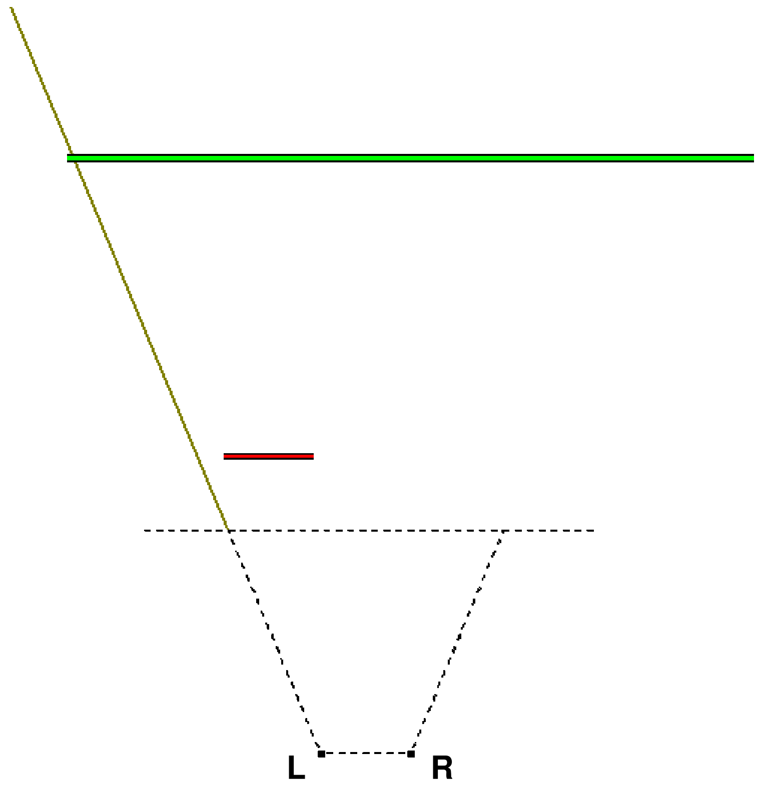

Construction

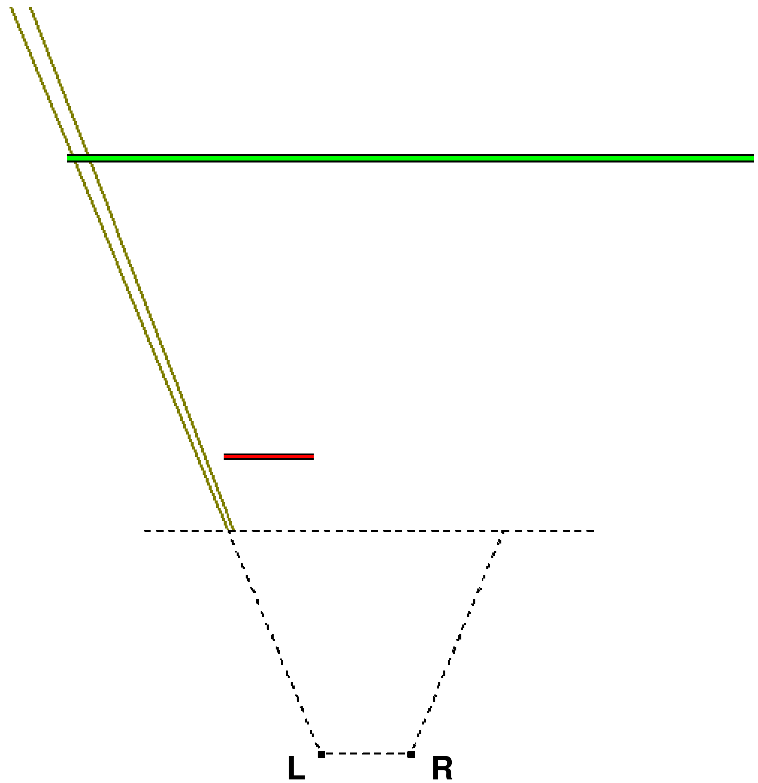

Construction

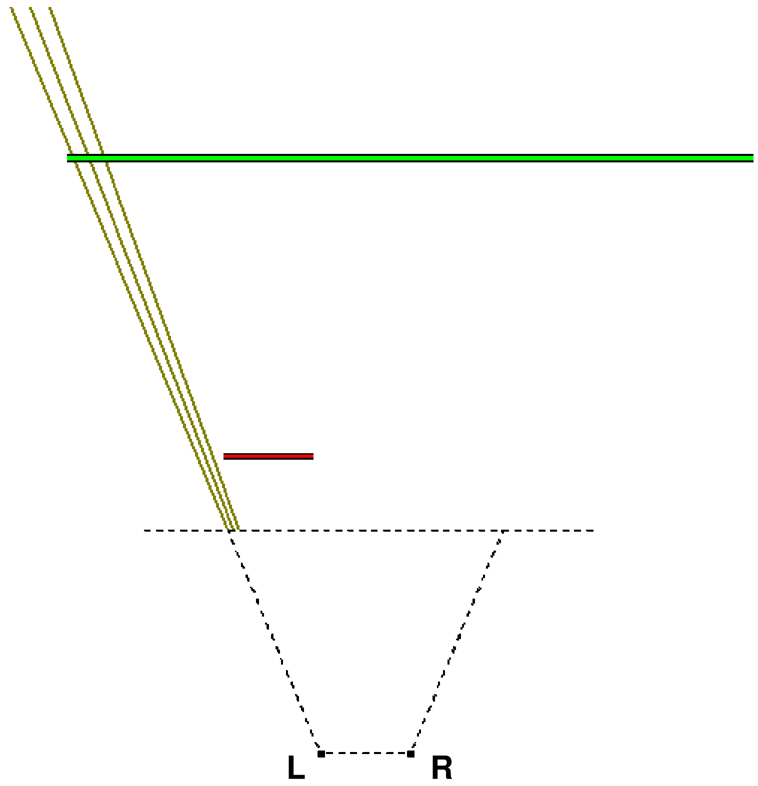

Construction

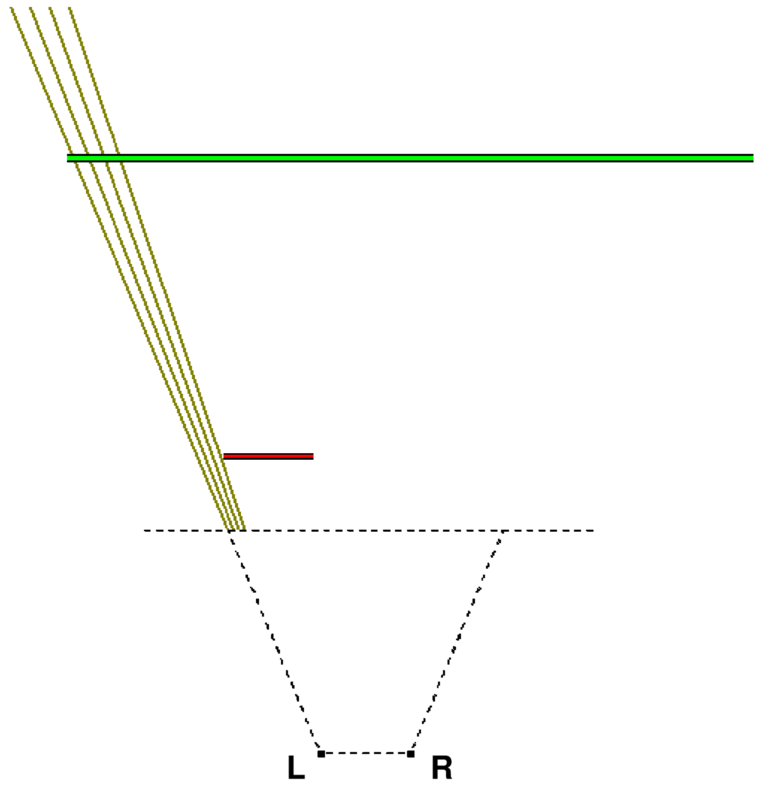

Construction

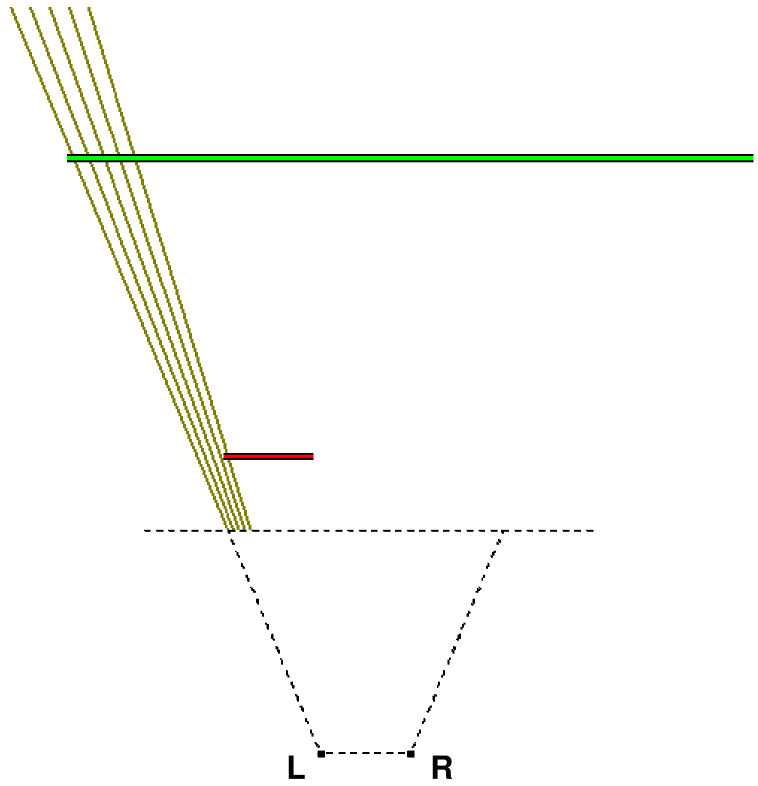

Construction

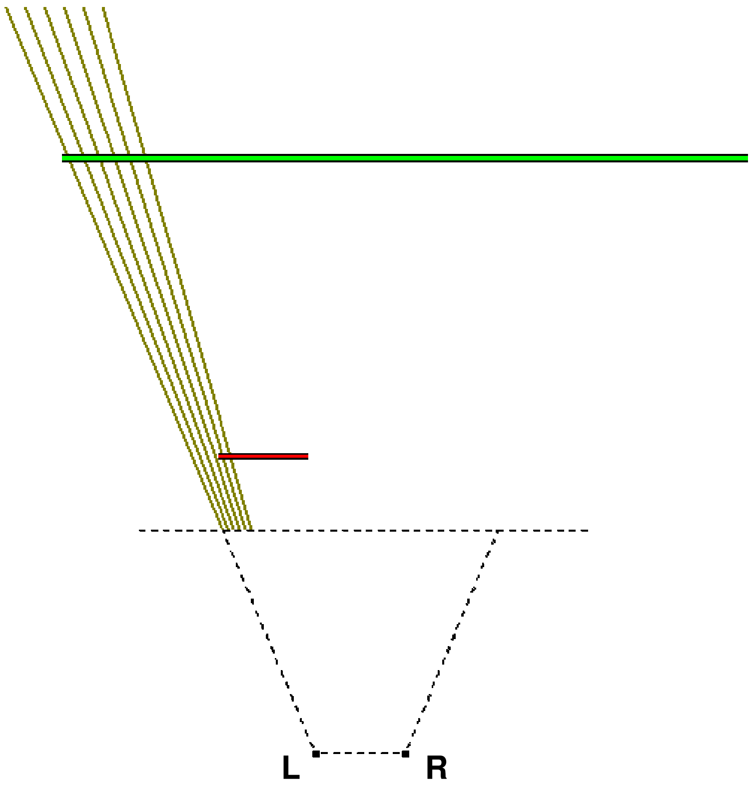

Construction

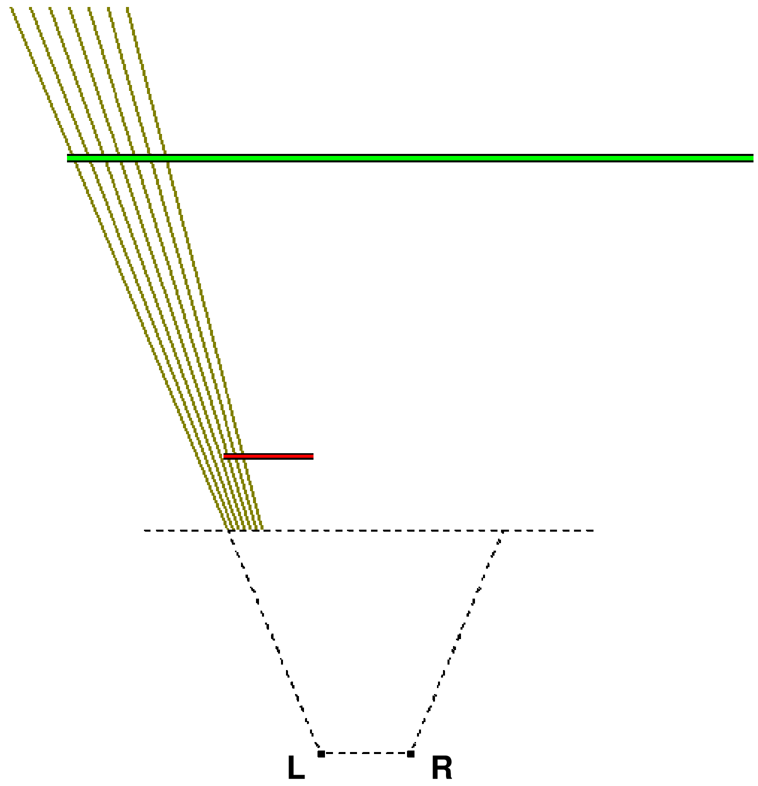

Construction

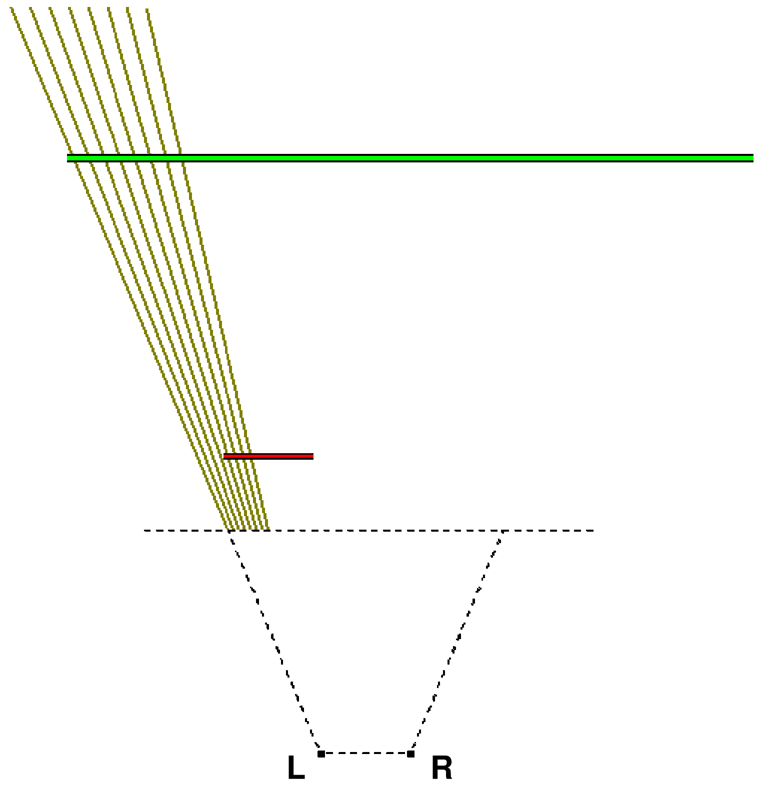

Construction

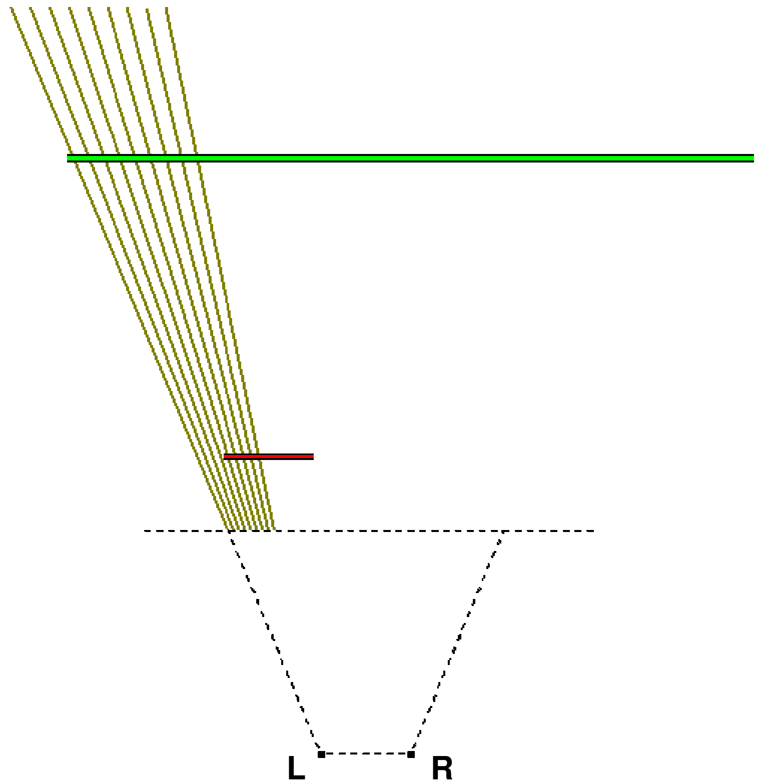

Construction

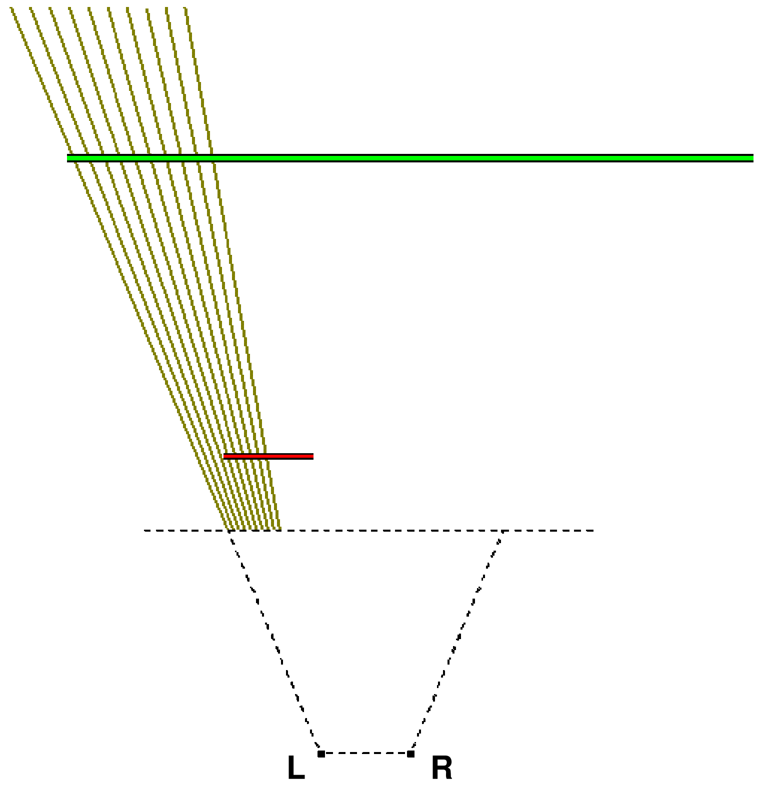

Construction

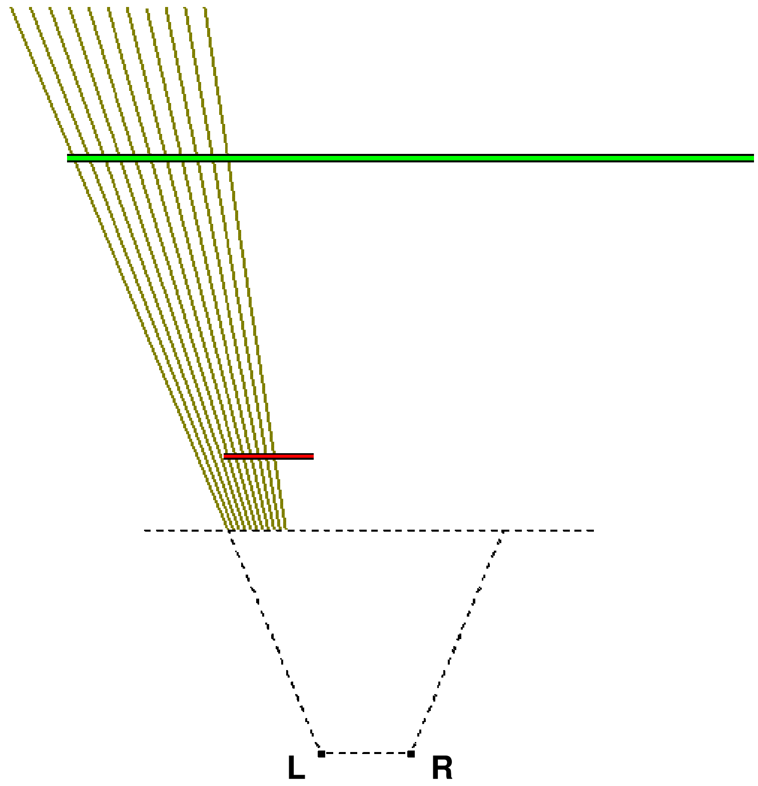

Construction

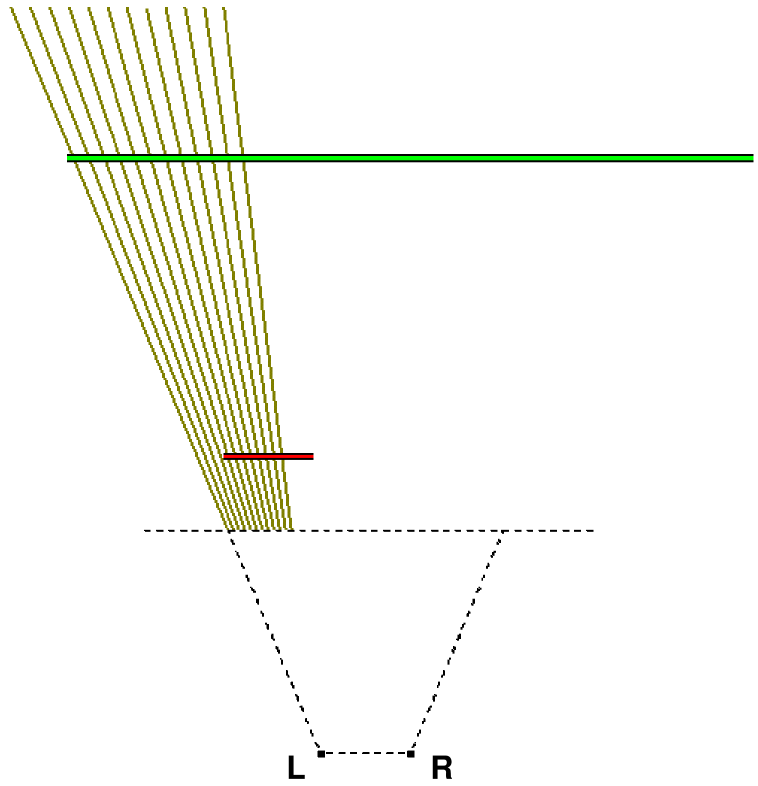

Construction

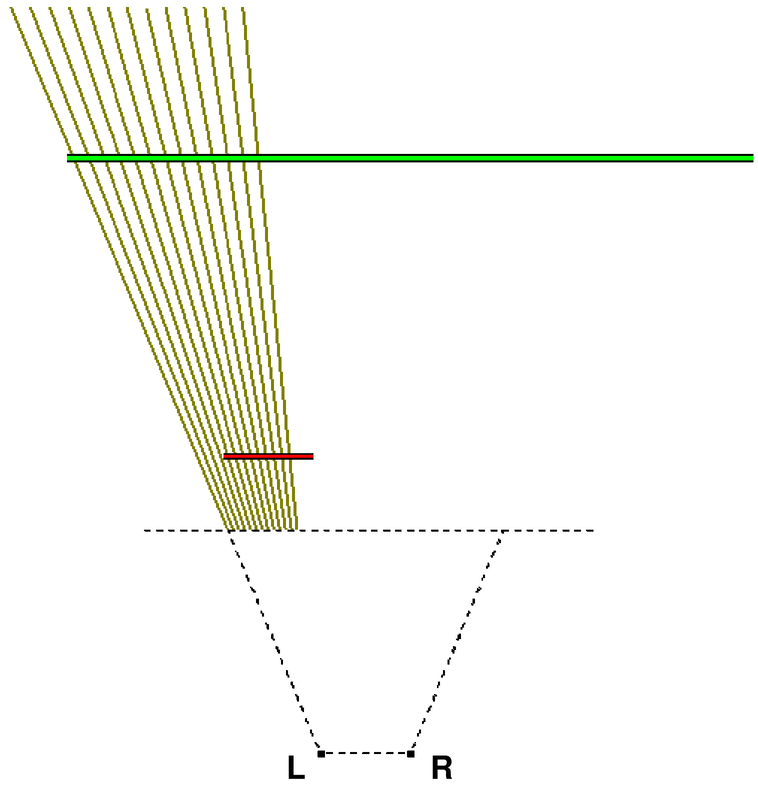

Construction

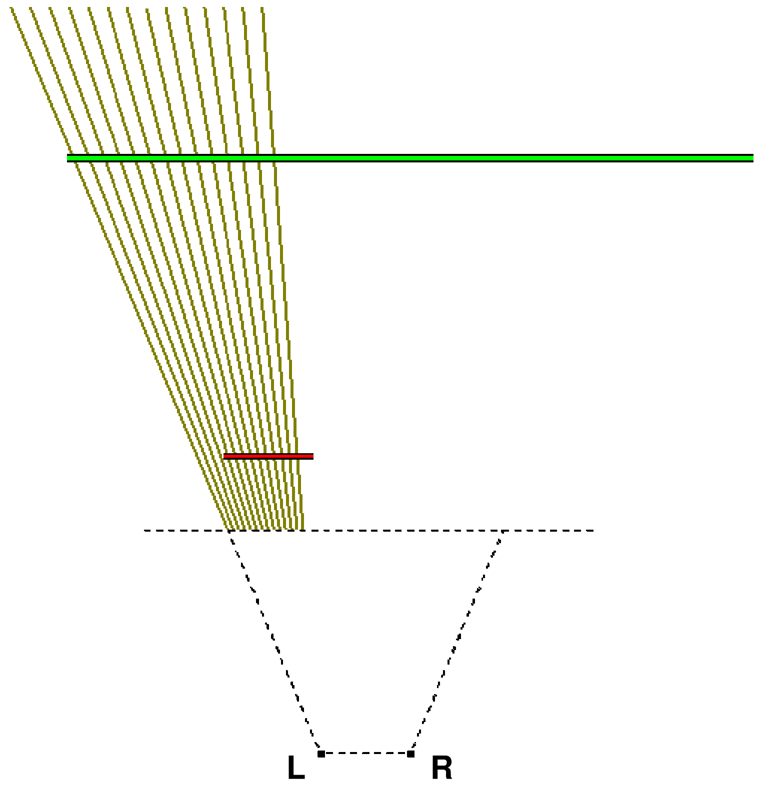

Construction

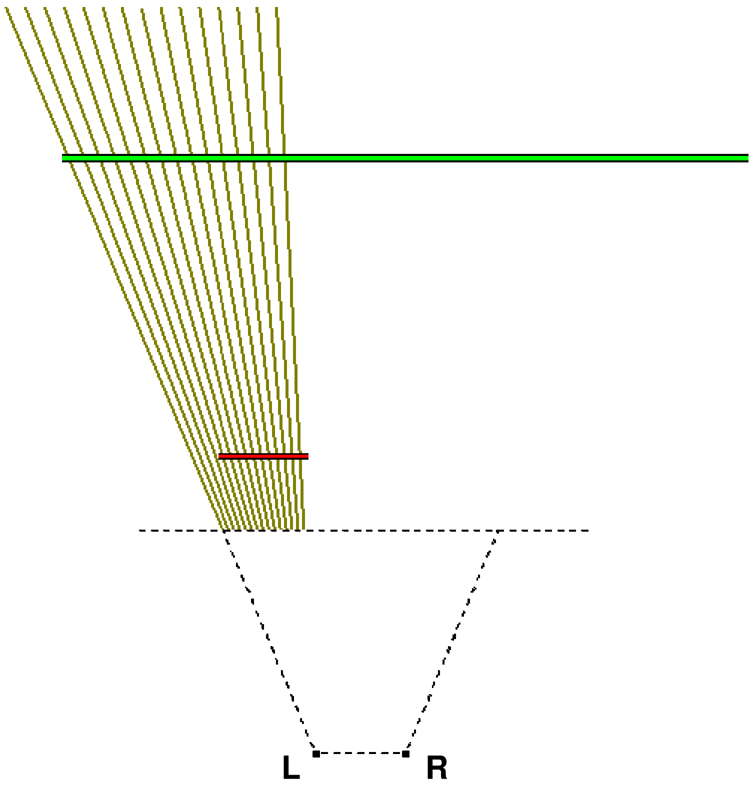

Construction

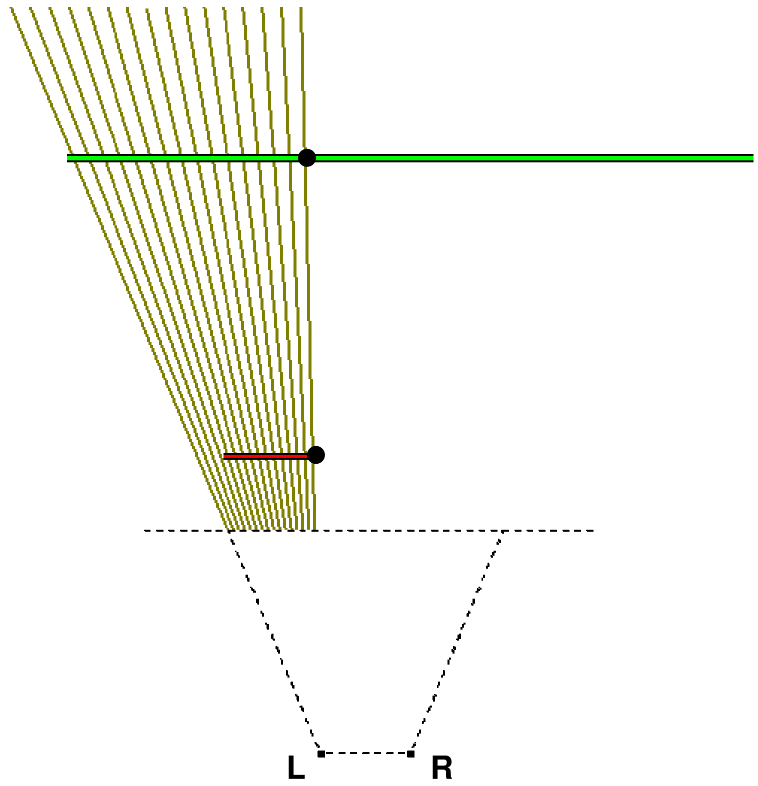

Construction

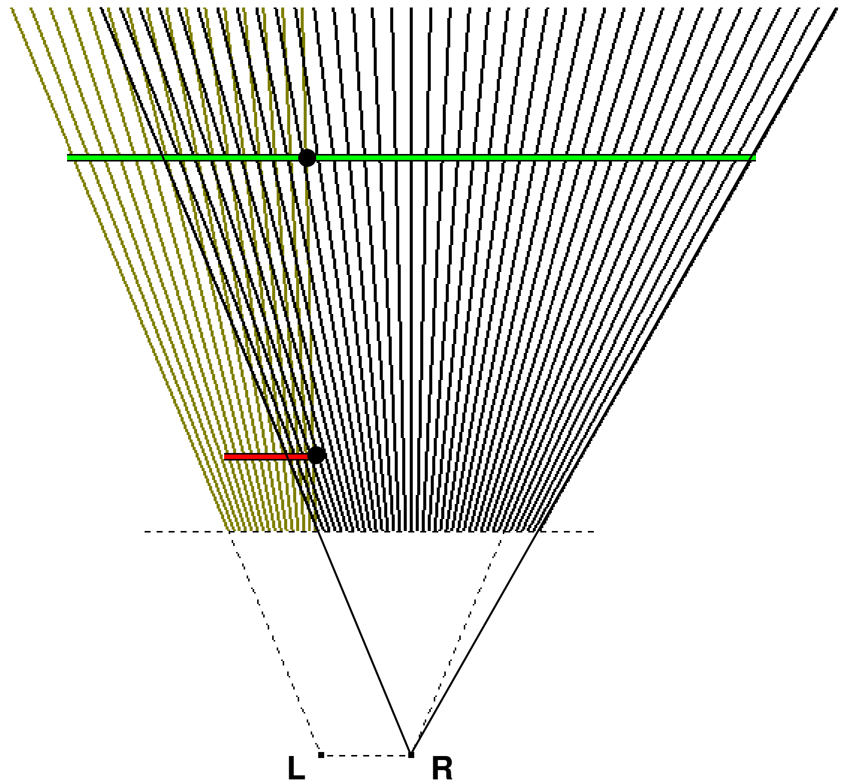

Construction

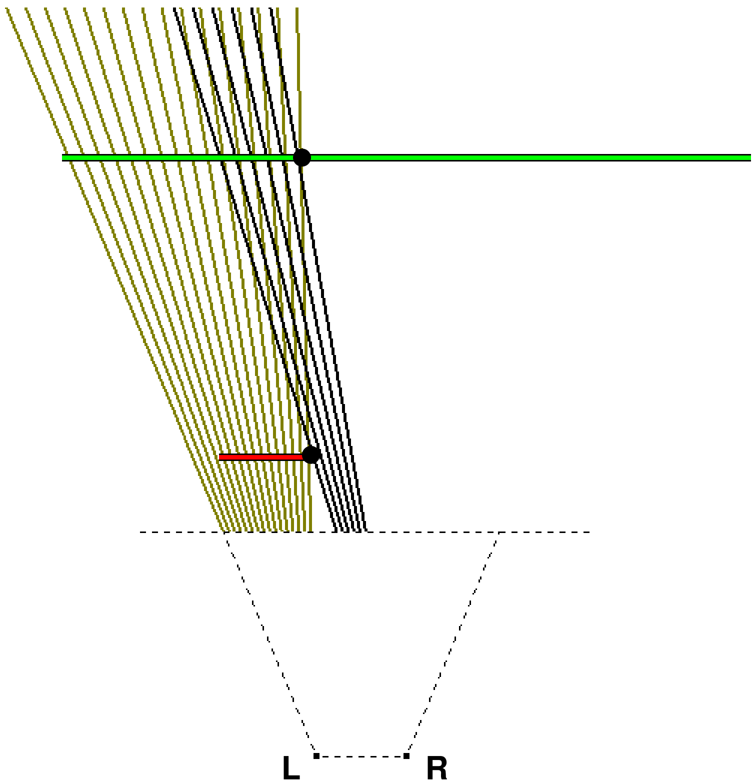

Construction

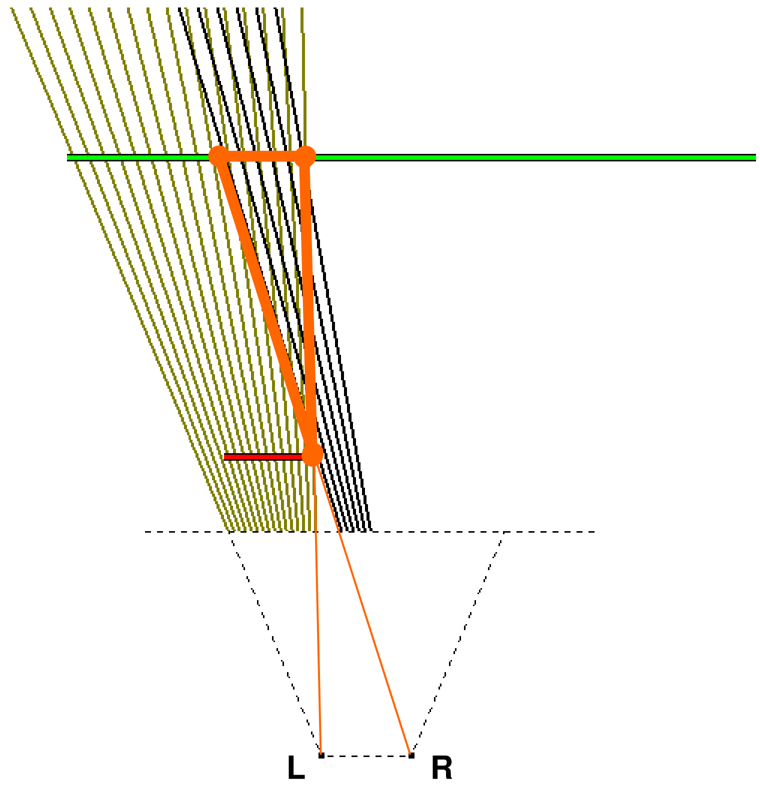

Construction

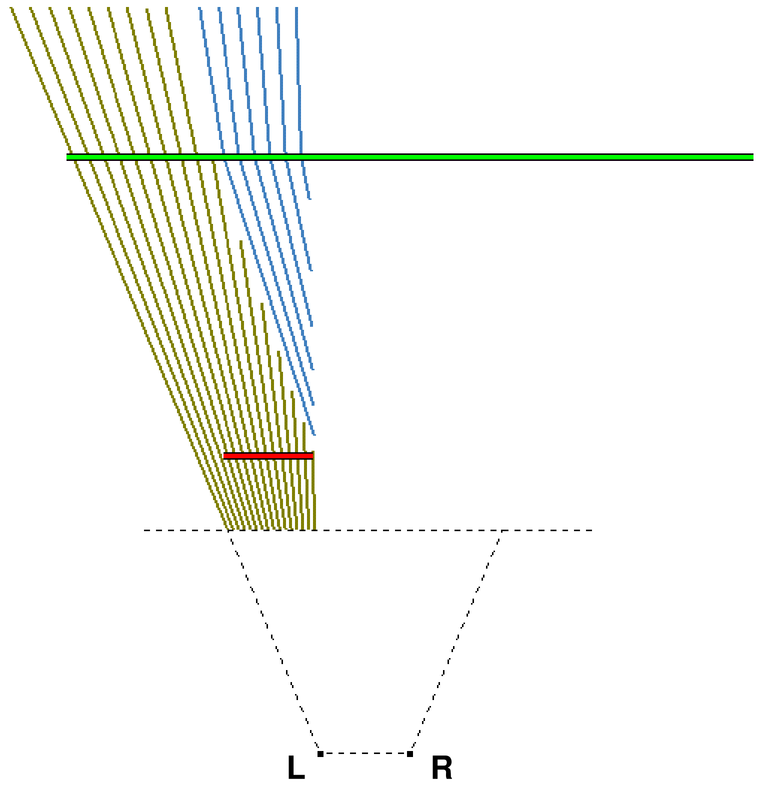

Construction

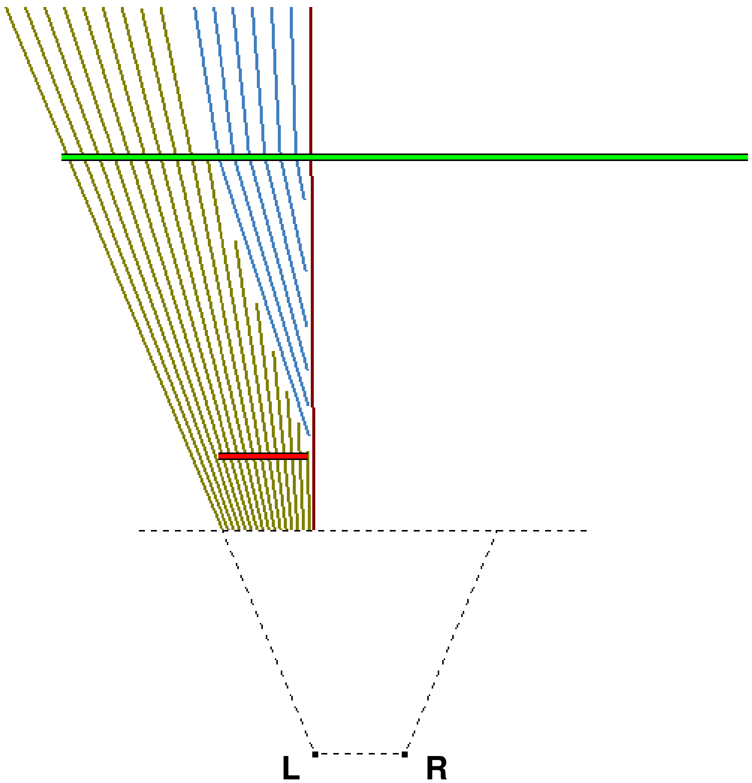

Construction

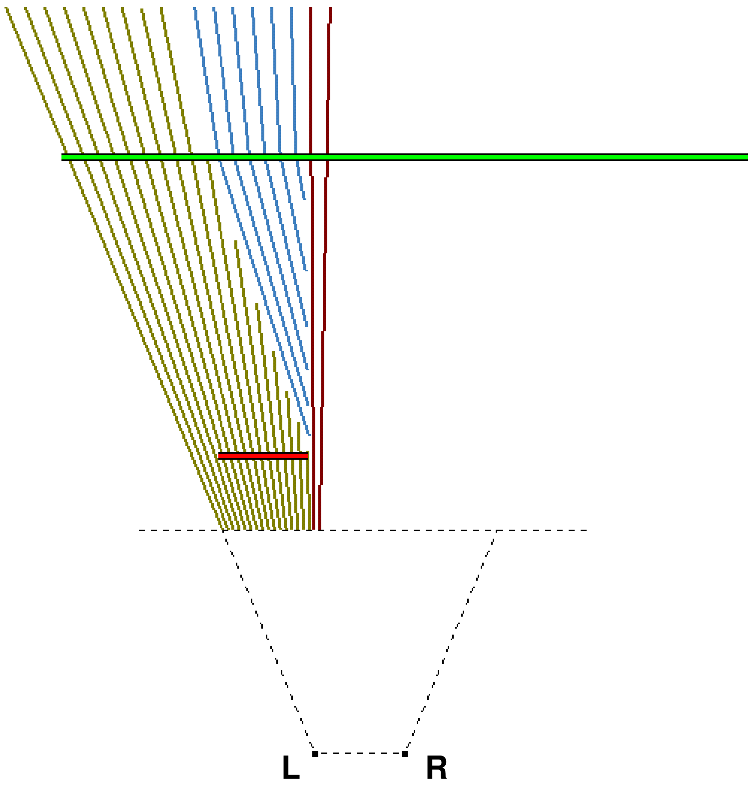

Construction

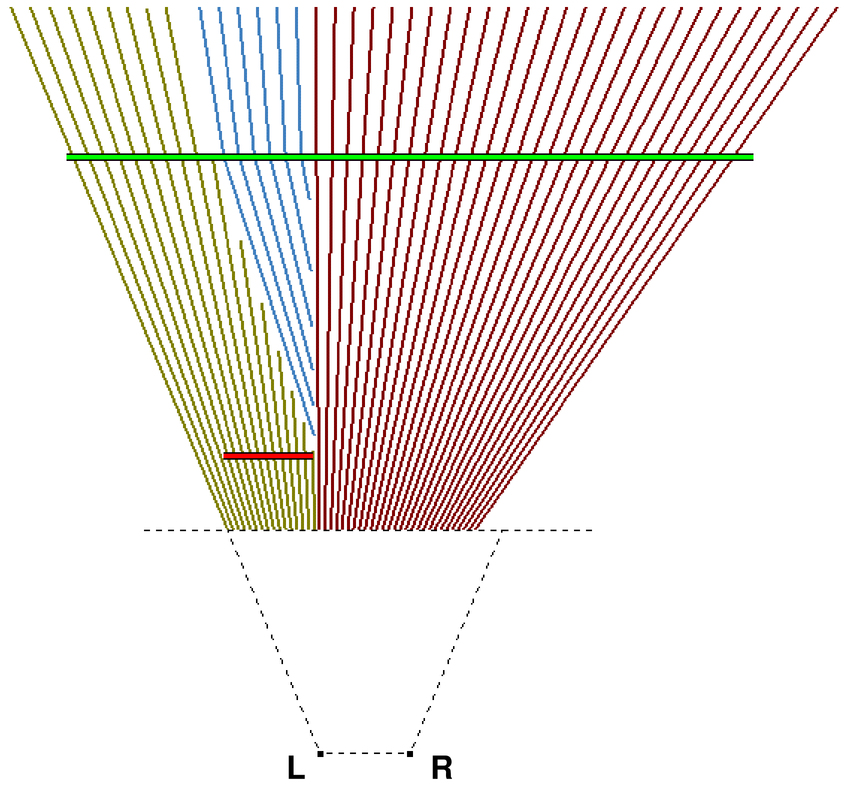

Construction

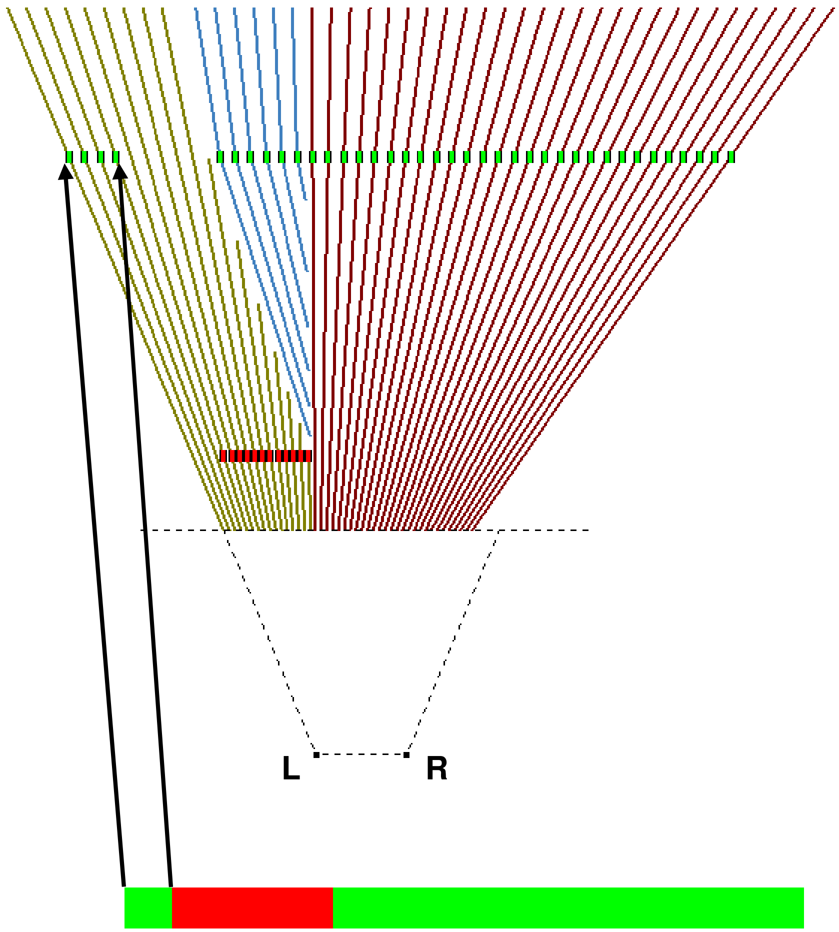

Construction

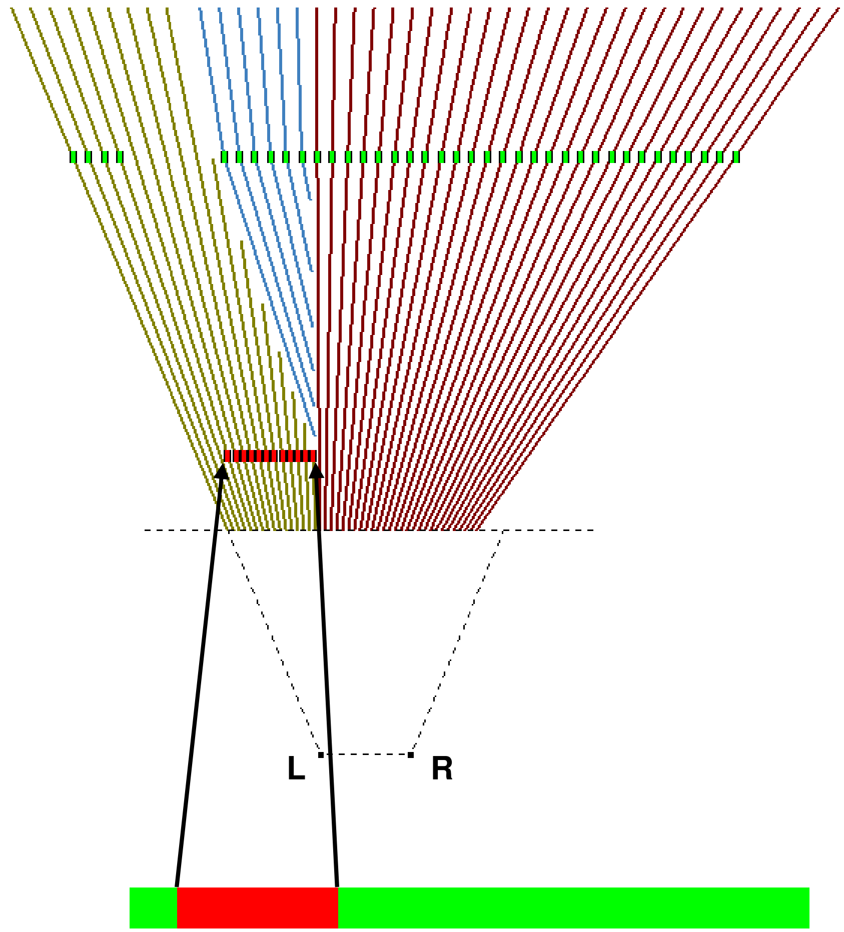

Construction

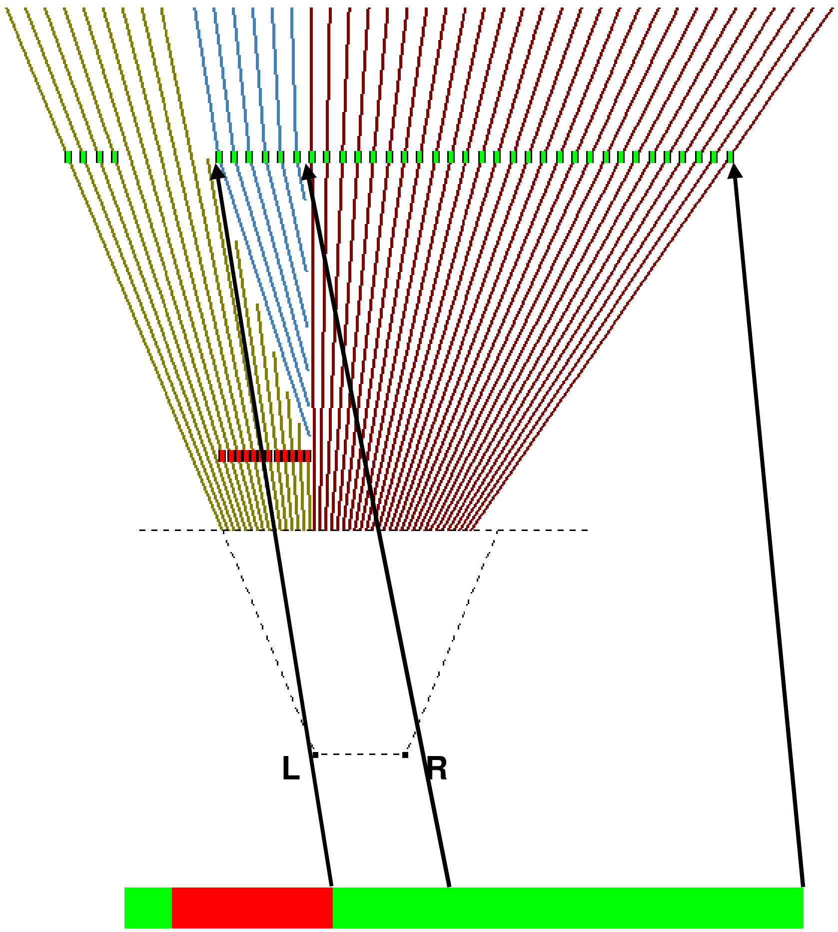

Construction
Thin Occluder

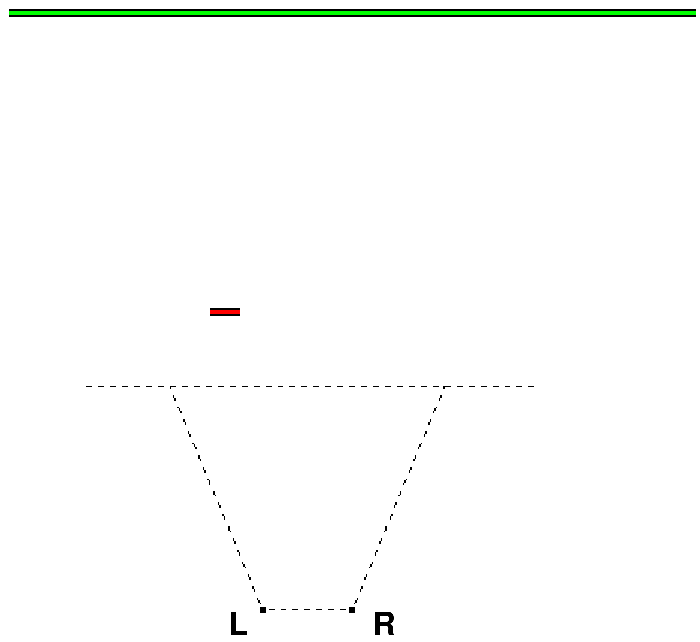

Construction
Thin Occluder

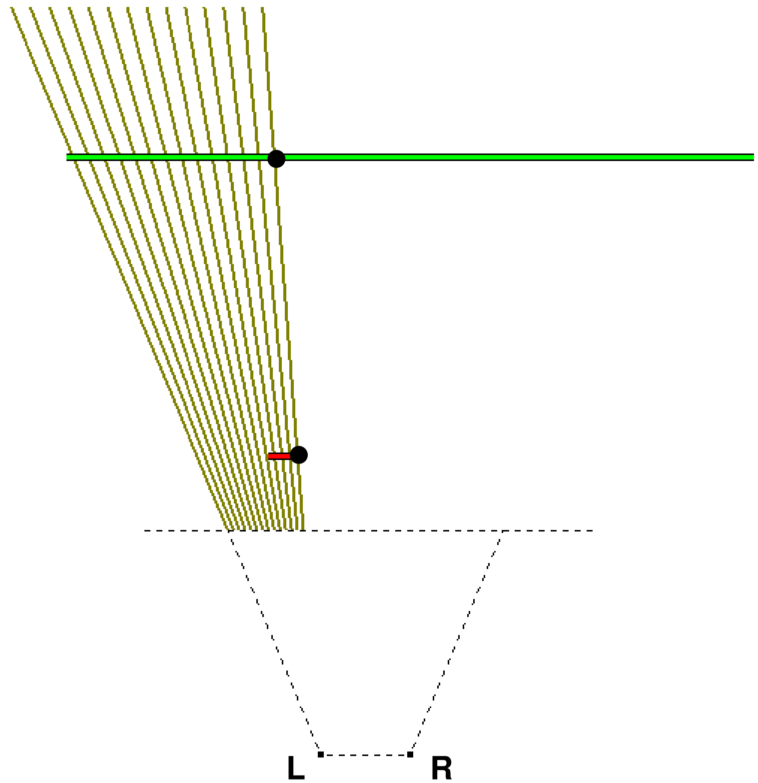

Construction
Thin Occluder

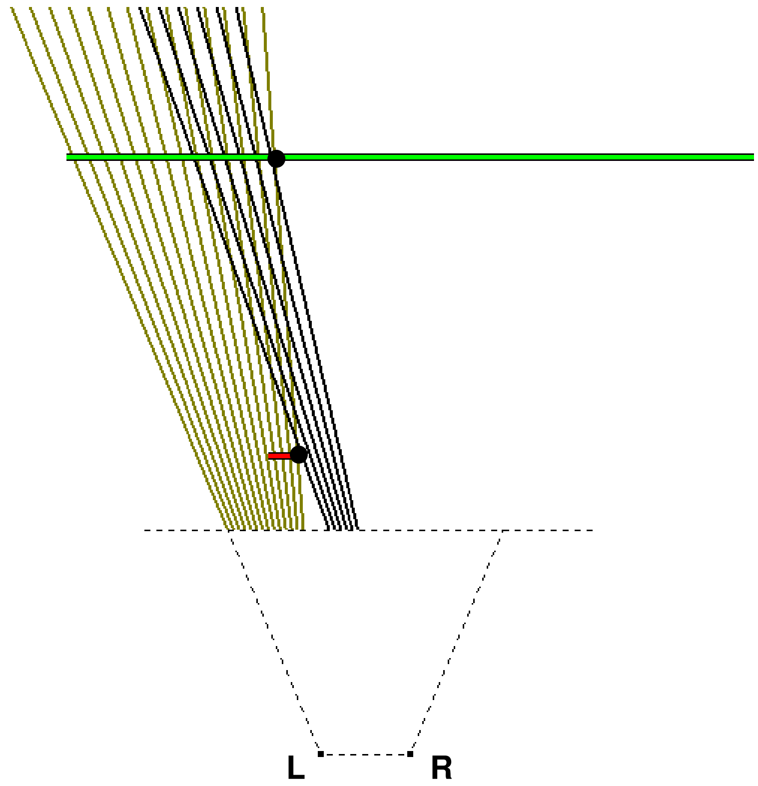

Construction
Thin Occluder

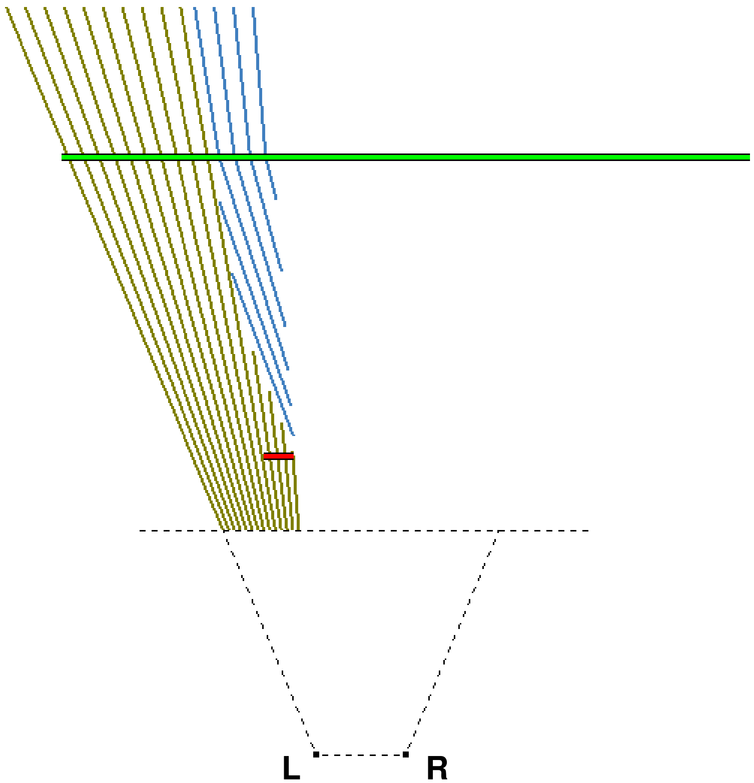

Construction
Thin Occluder

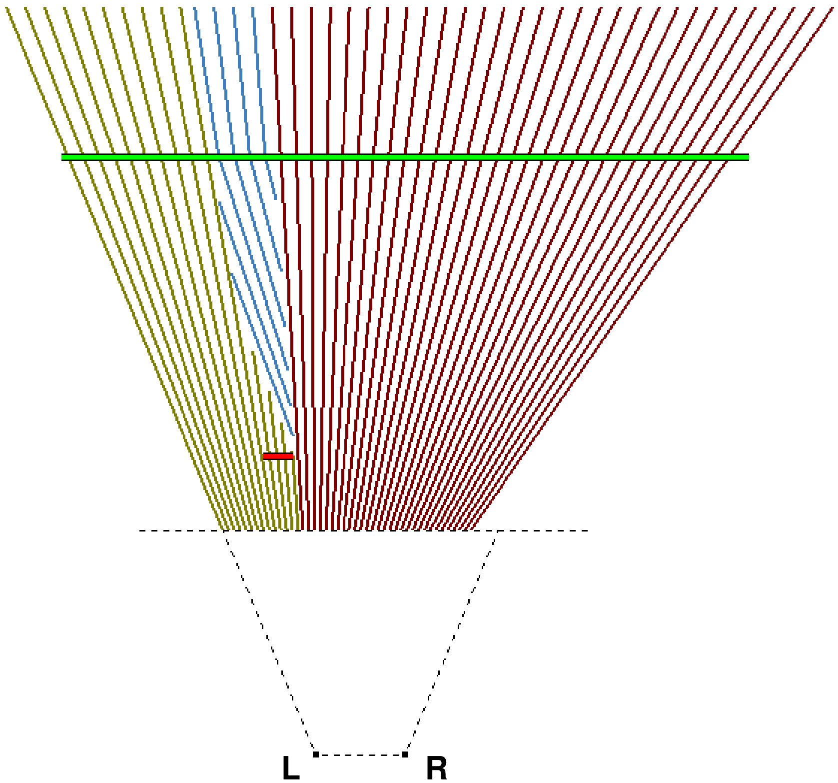

Construction
Thin Occluder

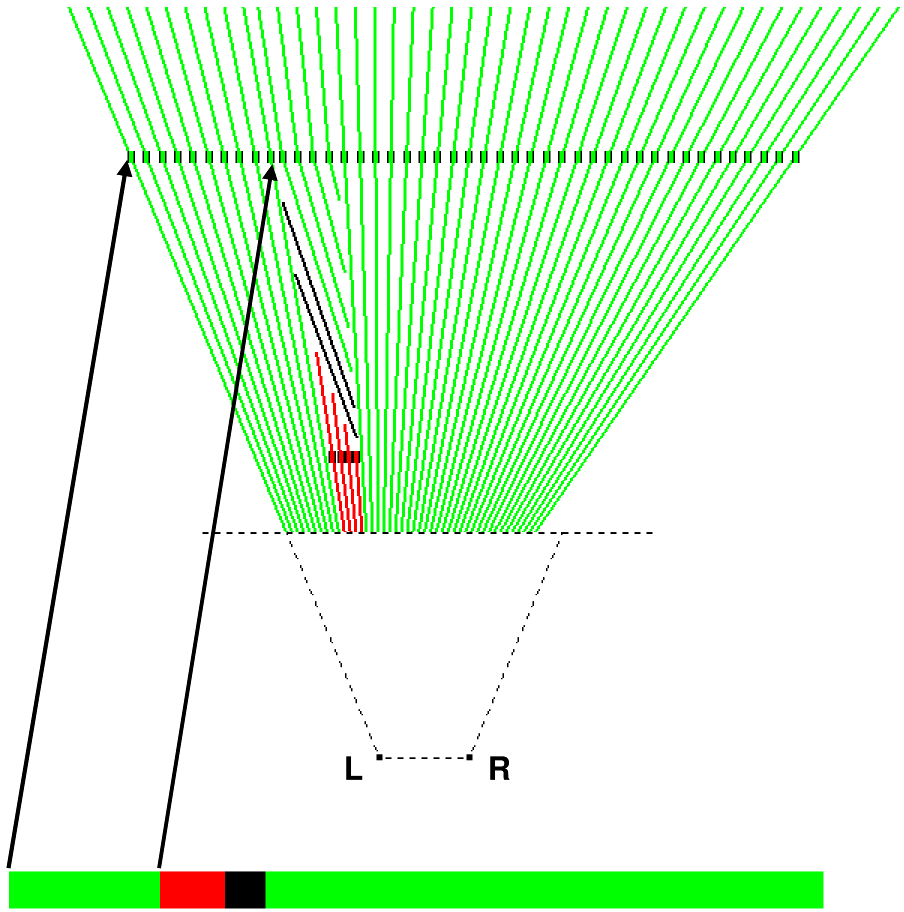

Construction
Thin Occluder

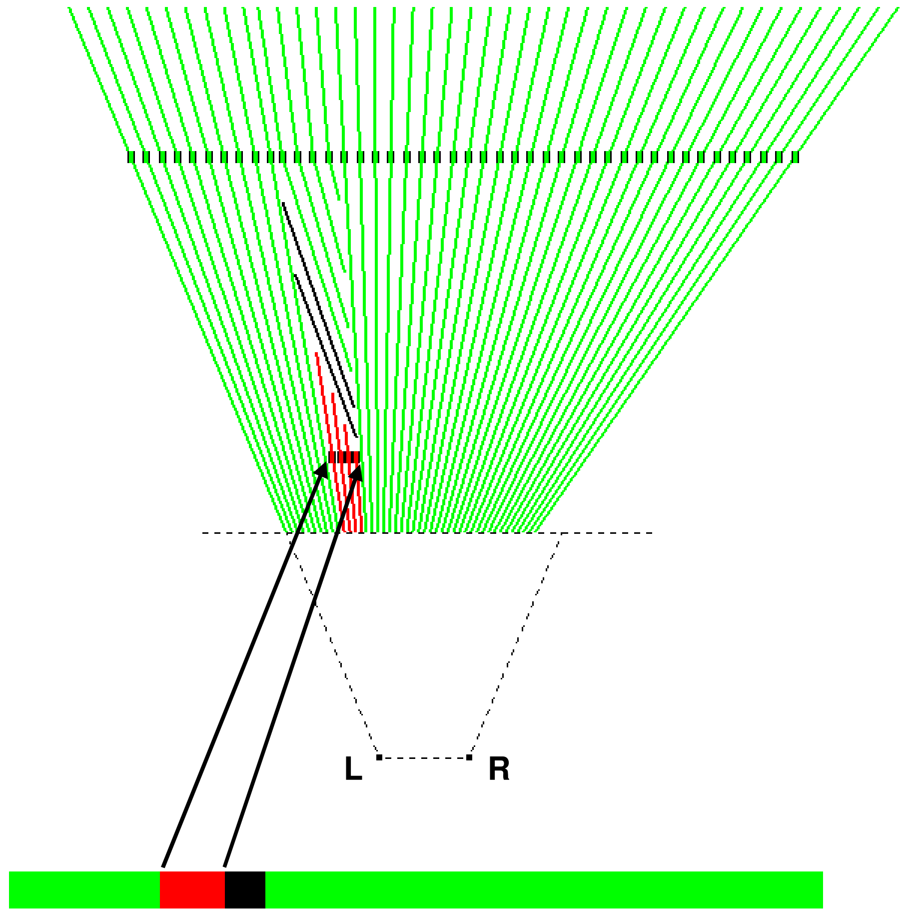

Construction Thin Occluder

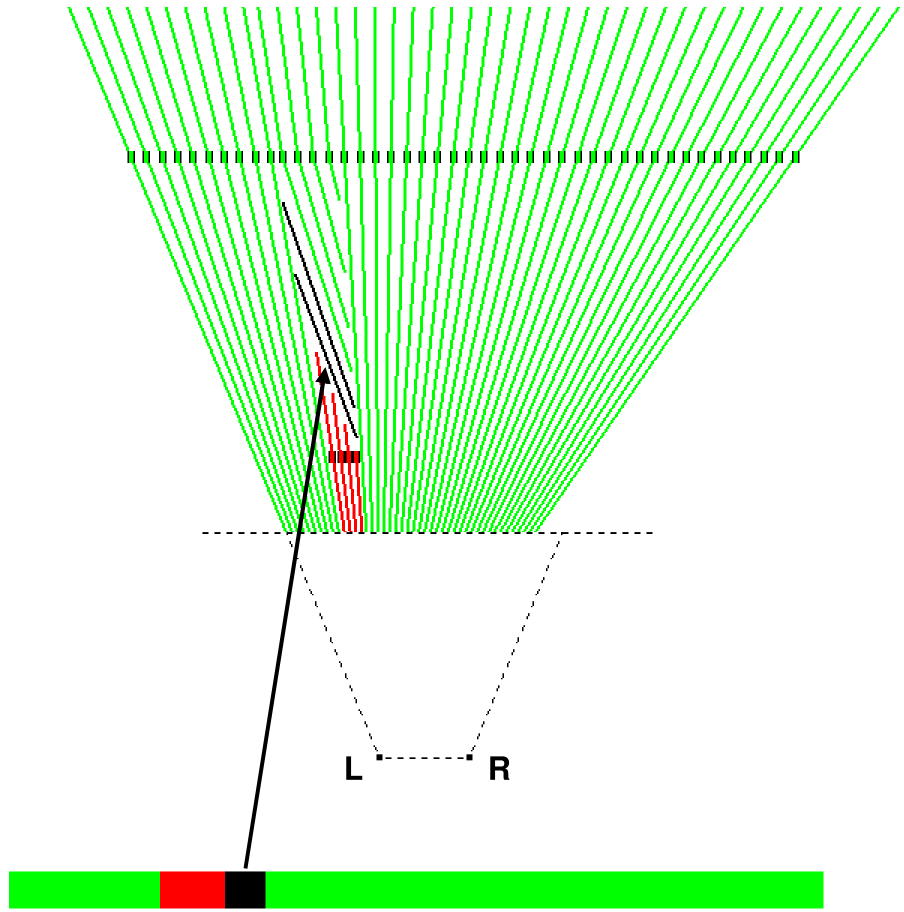

Construction
Thin Occluder

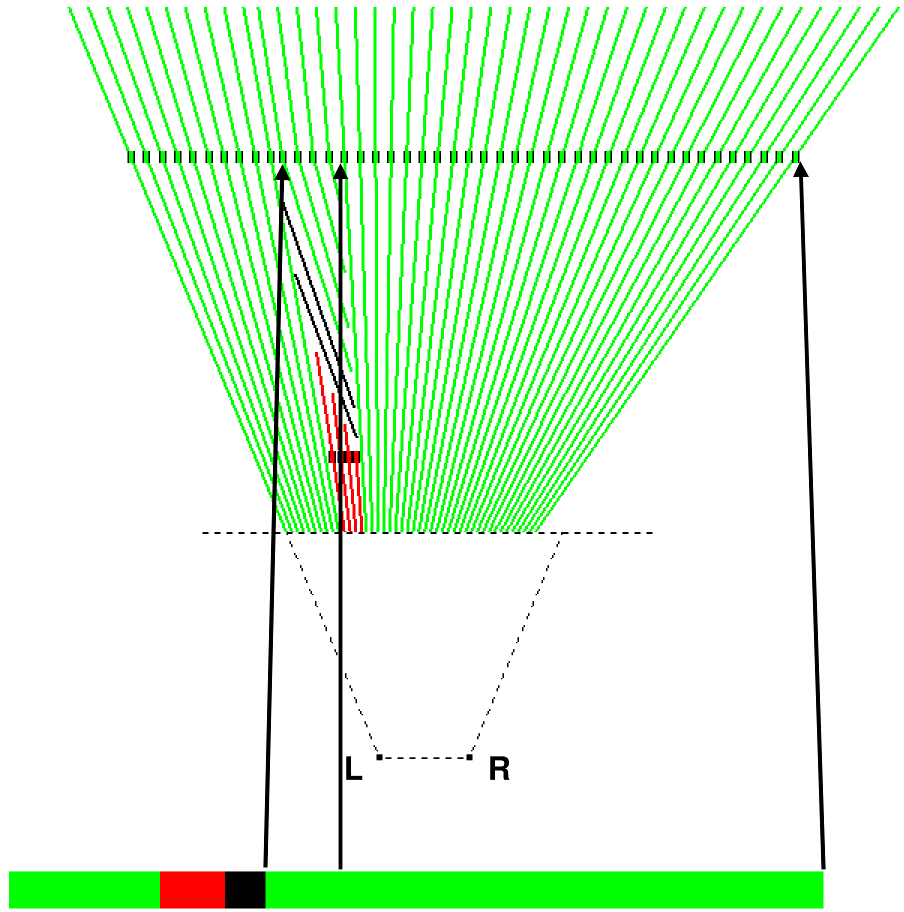

Construction examples

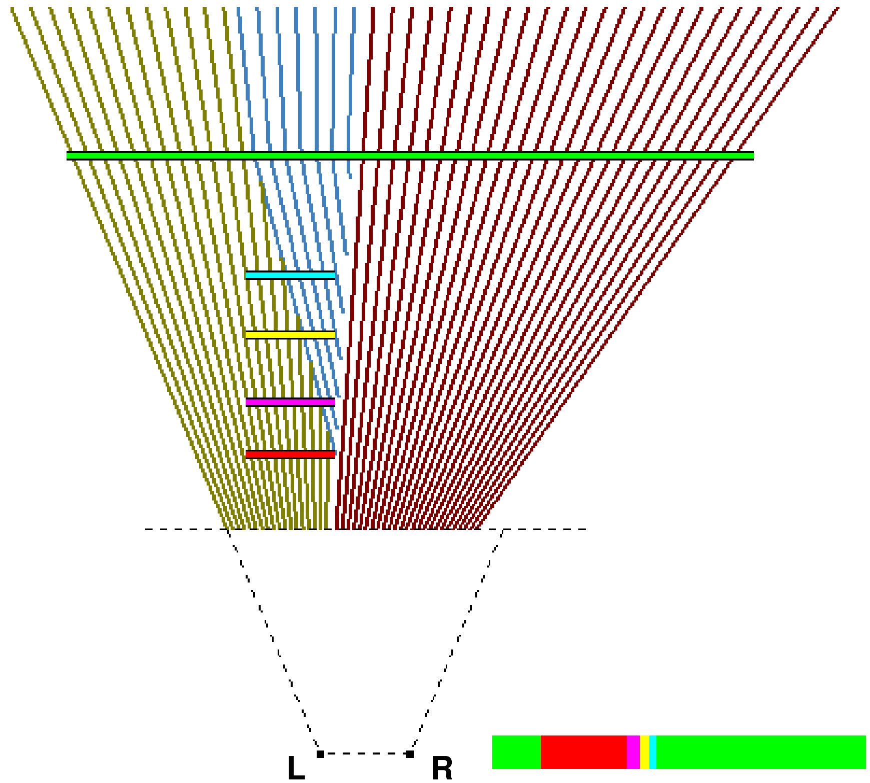

Construction examples

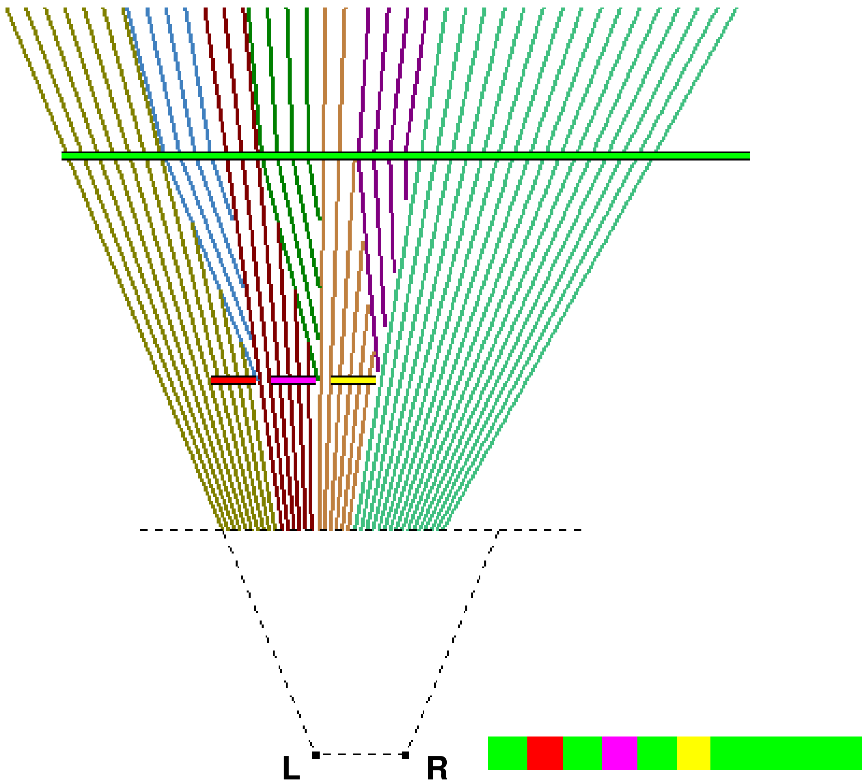

Rendering

- Given a camera model, one option is ray tracing
  - Trace rays back into 3-D scene
  - Compute intersections with geometric primitives
- Feed-forward EOC rendering algorithm on GPU
  - For each scene triangle
    - Project vertices (*vertex shader*)
    - Split into epipolar spans (*geometry shader*)
    - For each span
      - Project span endpoints (*geometry shader*)
      - Rasterize projected span (*pixel shader*)

Rendering

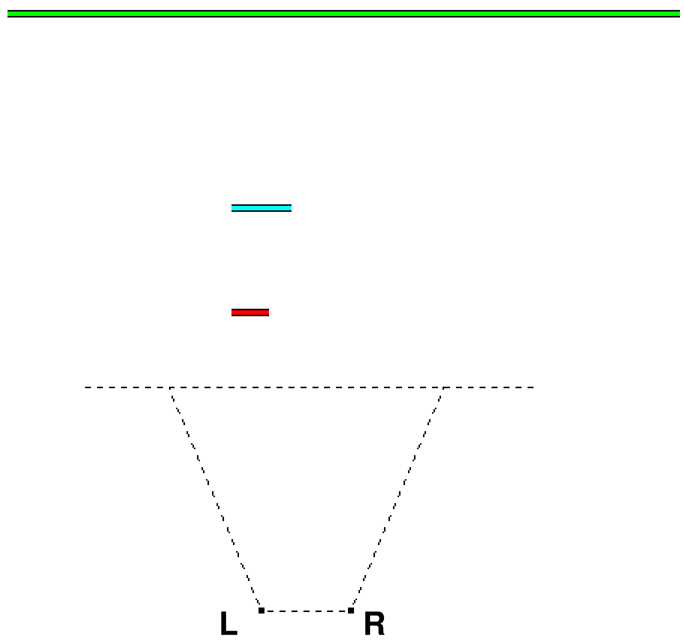

Rendering

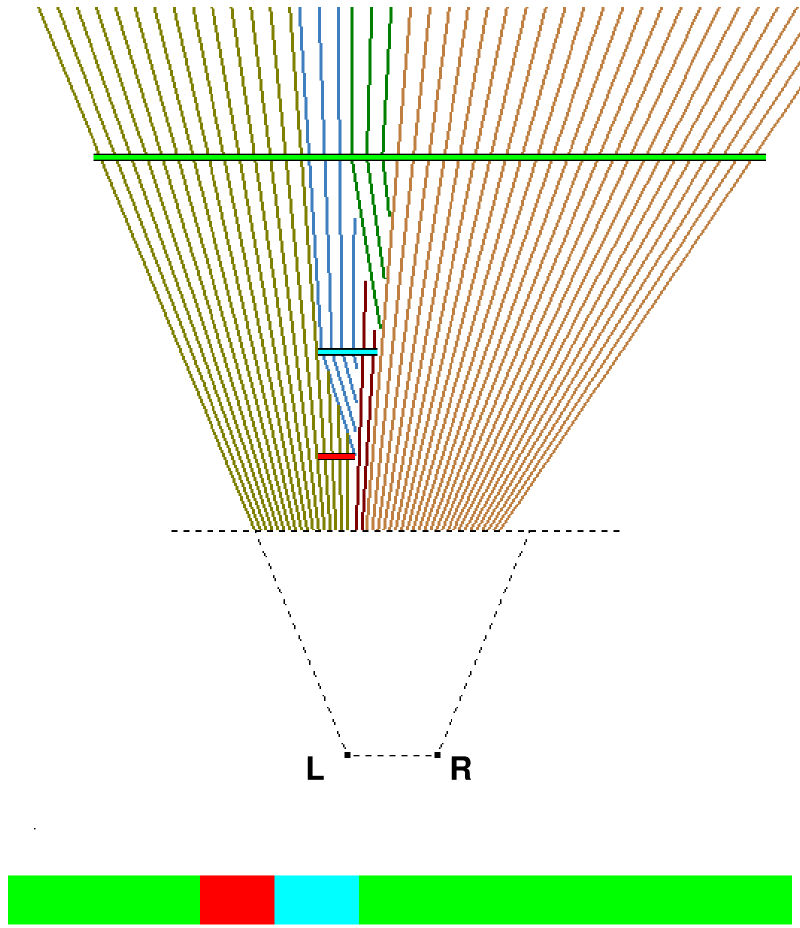

Rendering

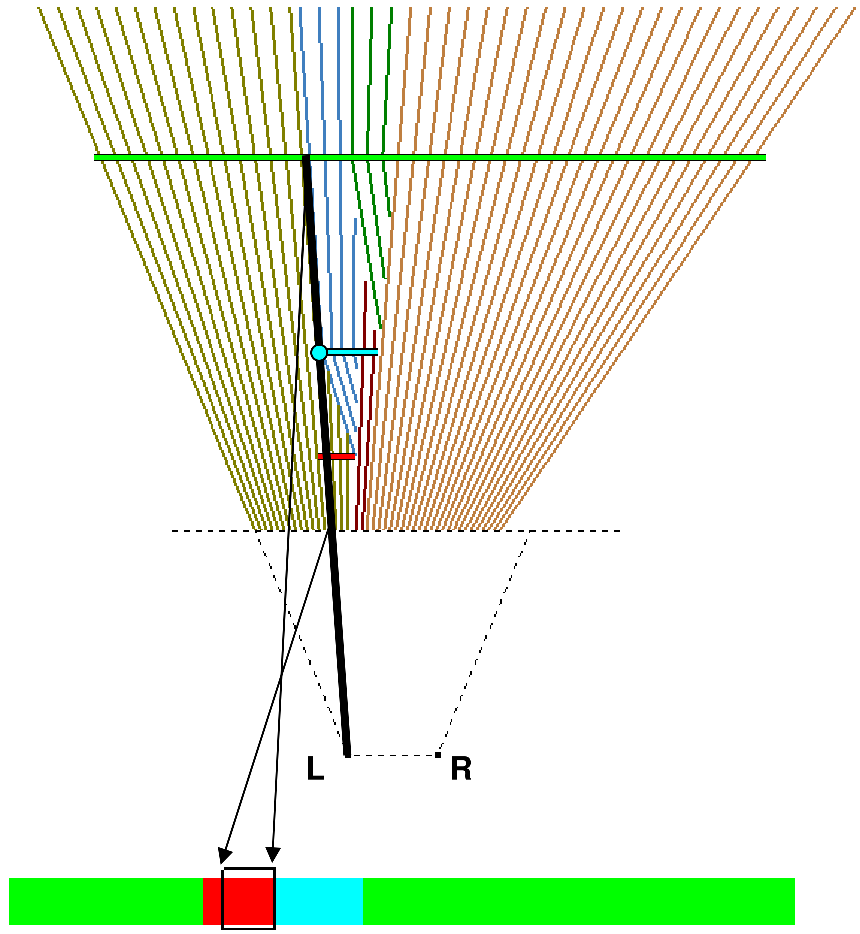

Rendering

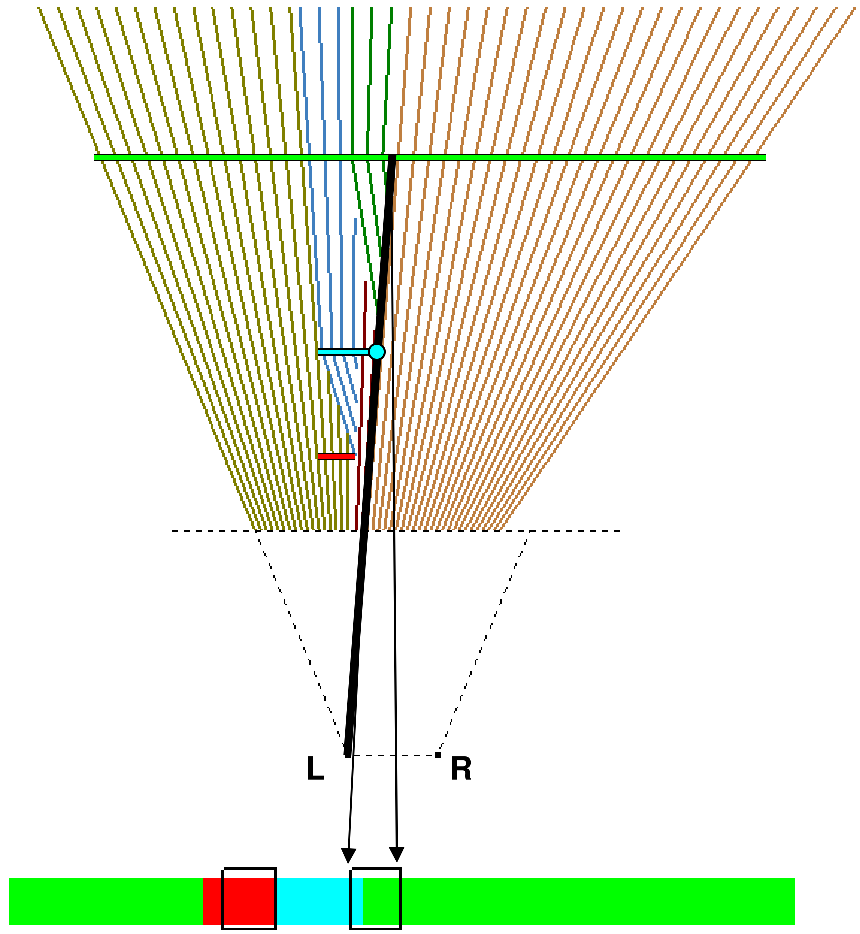

Rendering

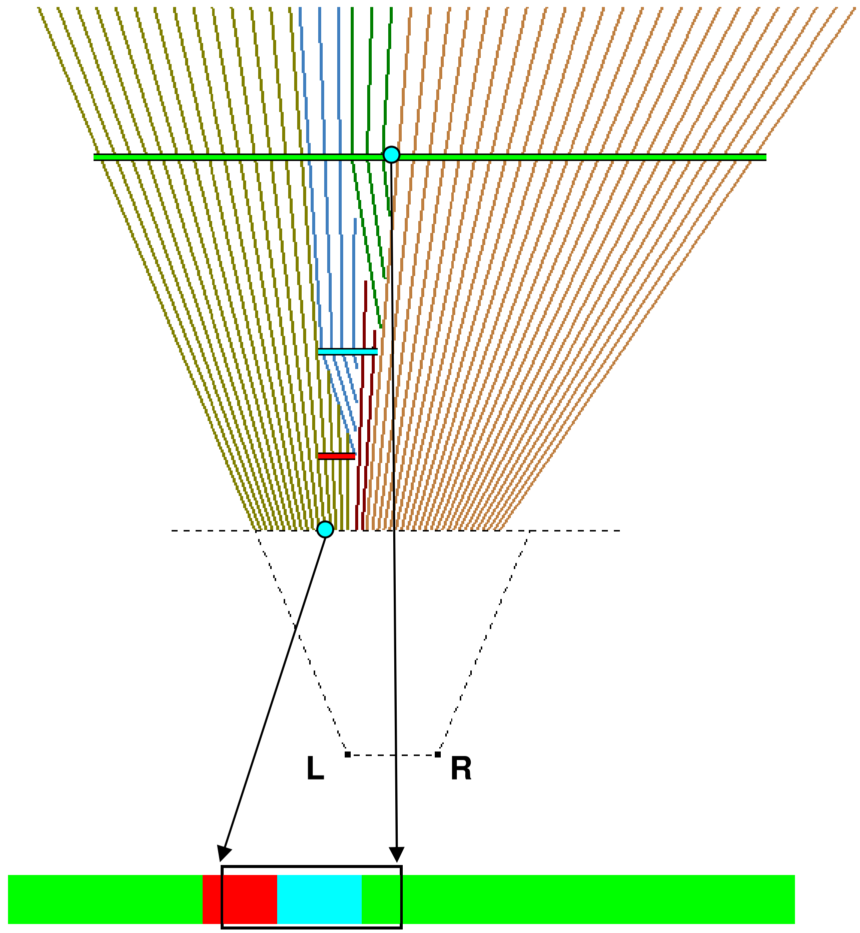

Rendering

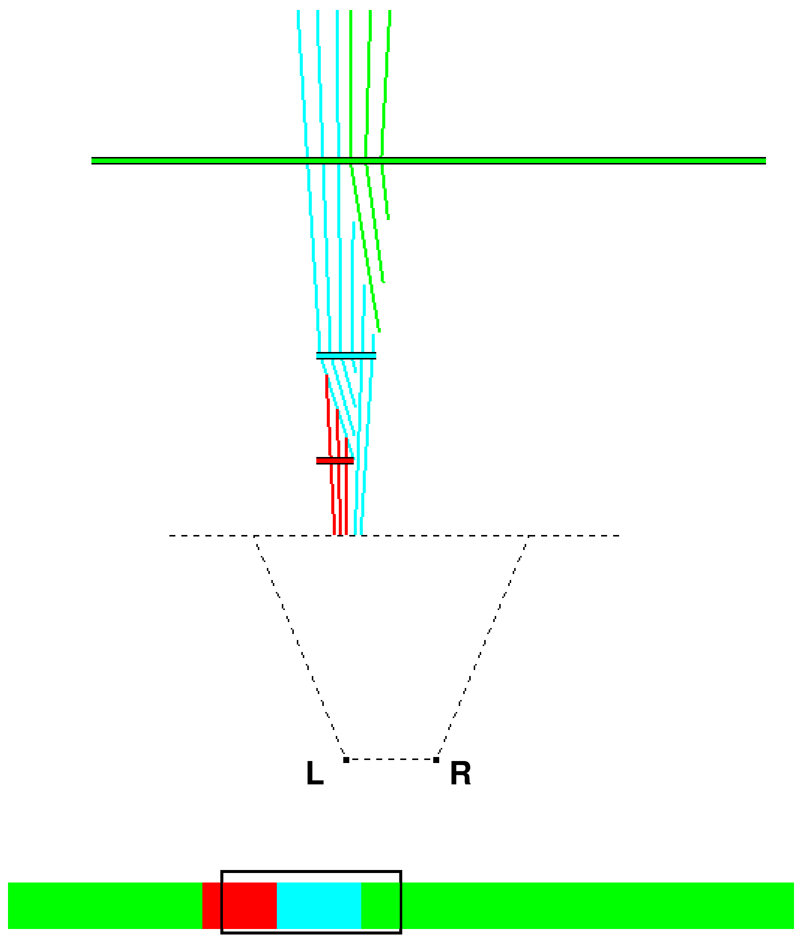

Picket Fence

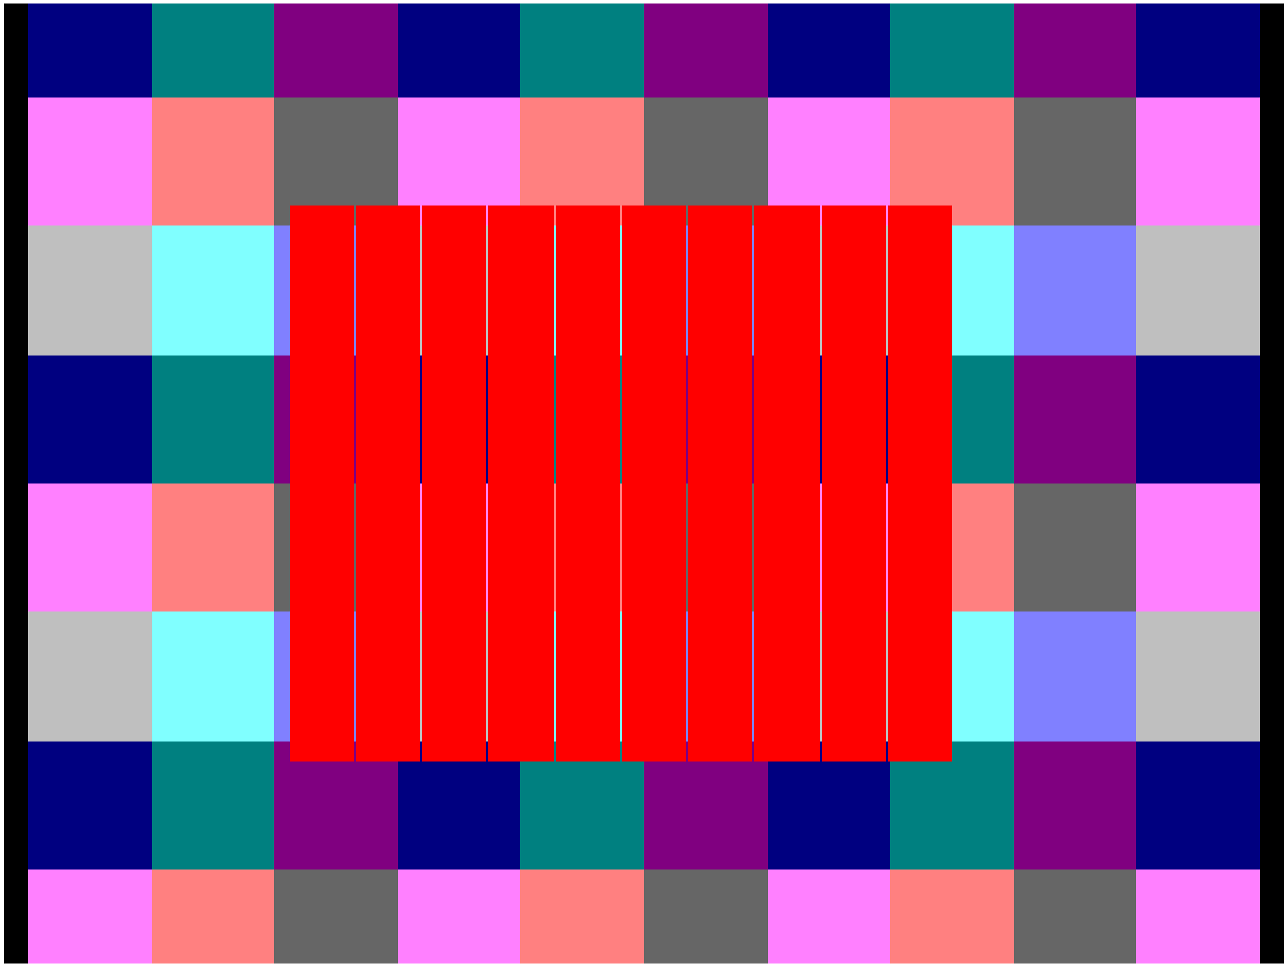

Picket Fence

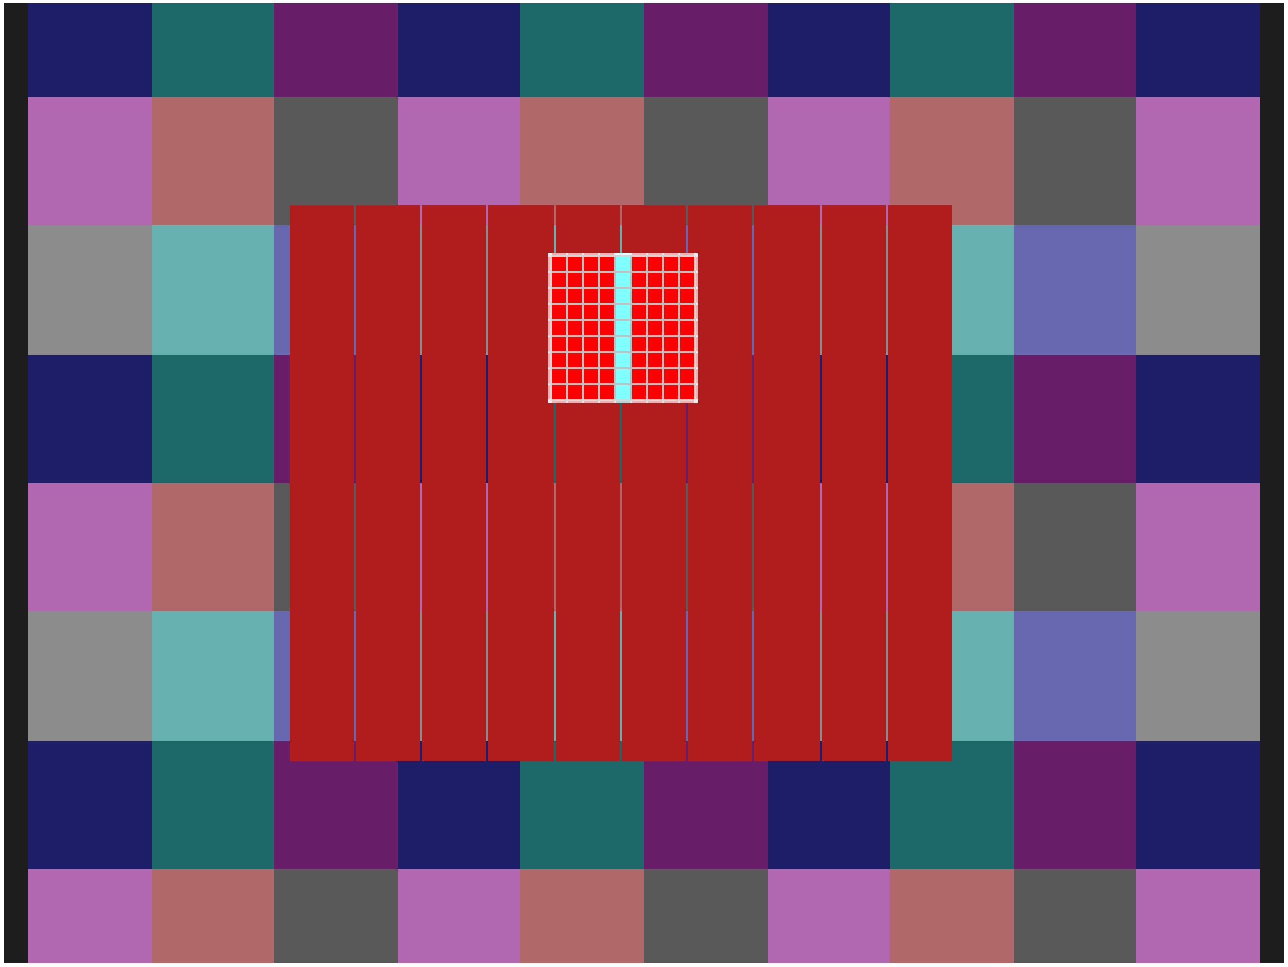

Picket Fence

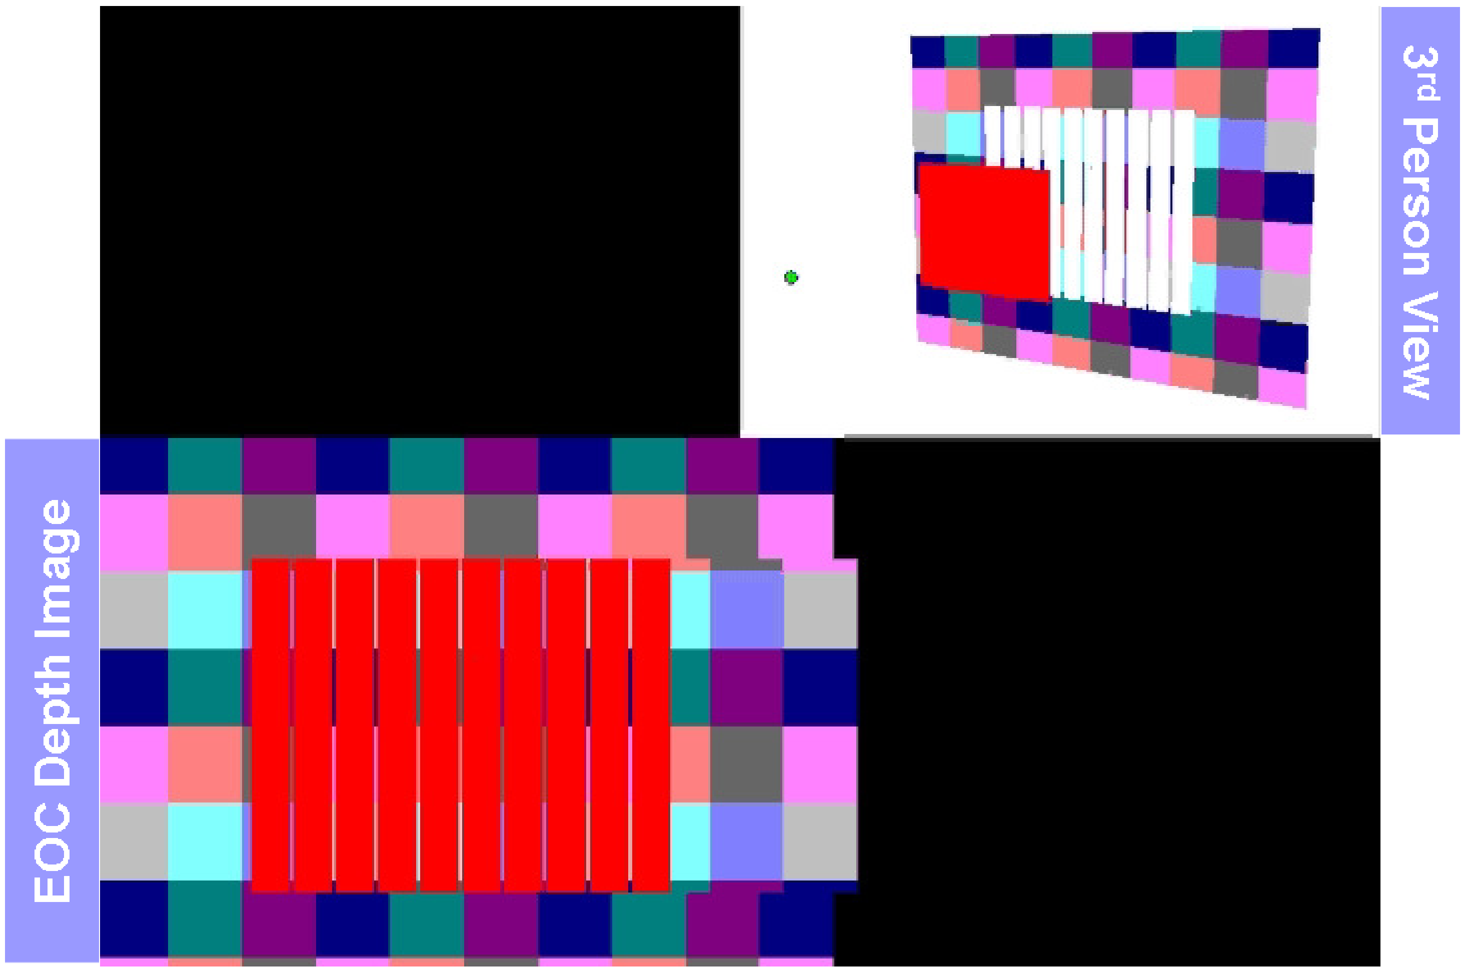

Using EOC Images

- Recovering 3-D positions is trivial
  - Rays are calculated during construction
  - Depth is calculated during rendering
- Scene Reconstruction
  - Point based rendering is trivial
  - Triangulation is more challenging

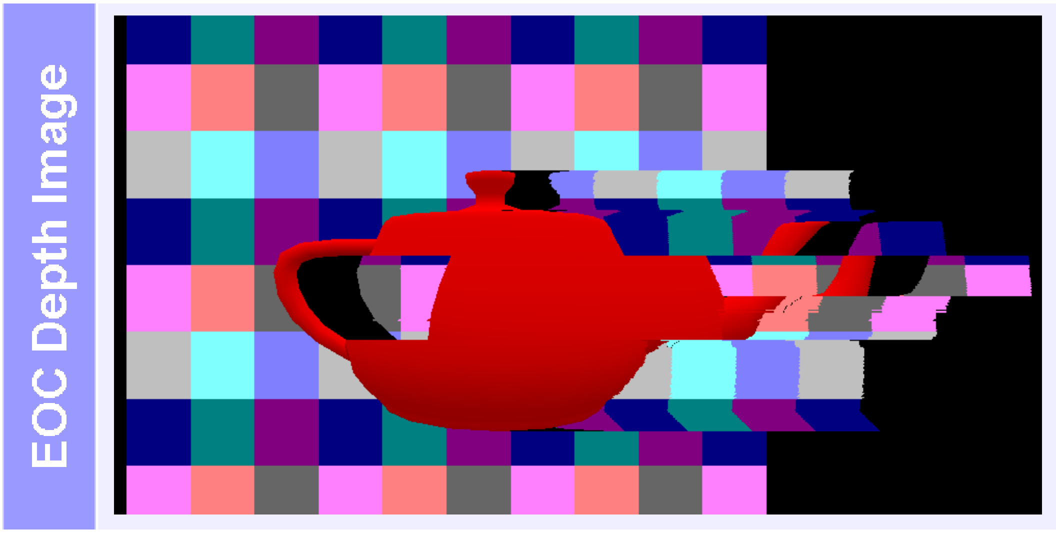

Scene Reconstruction

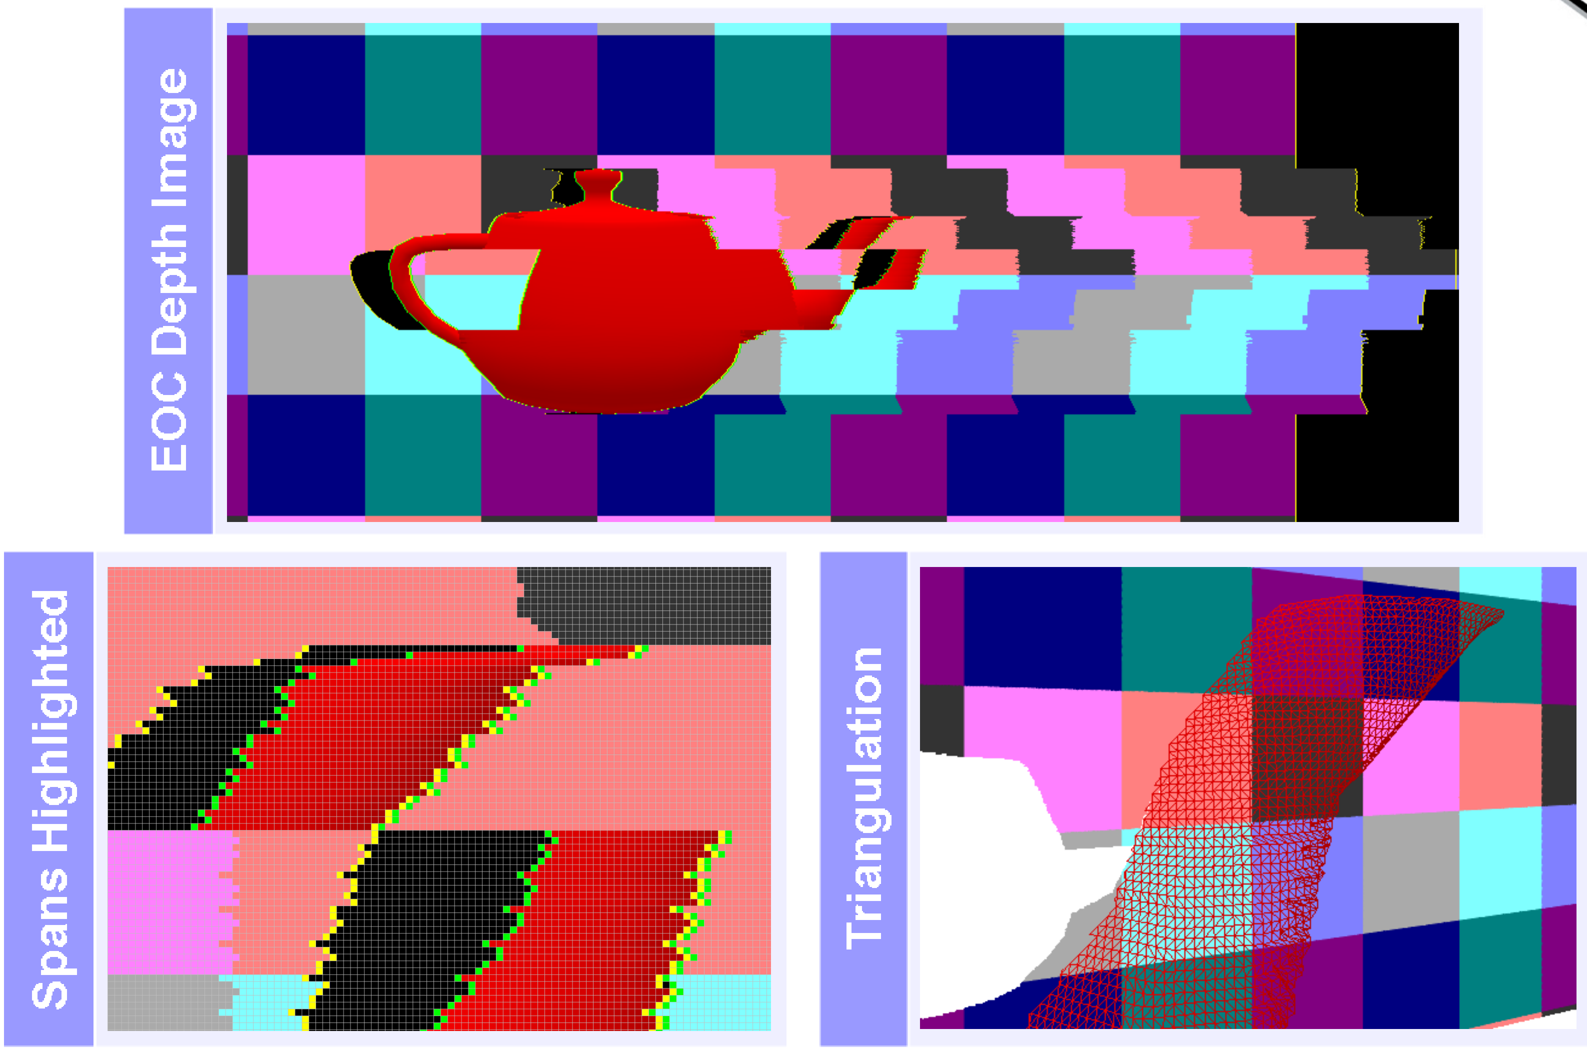

Radial Epipolar Motion Sphere

PHC Projection	Radial Projection
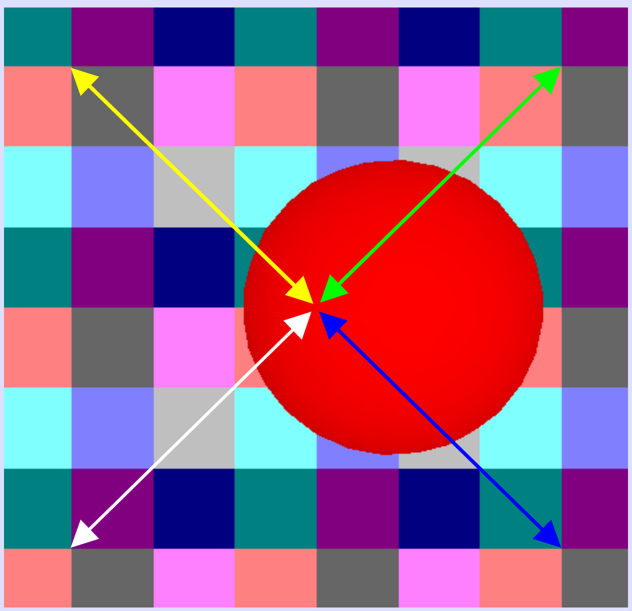	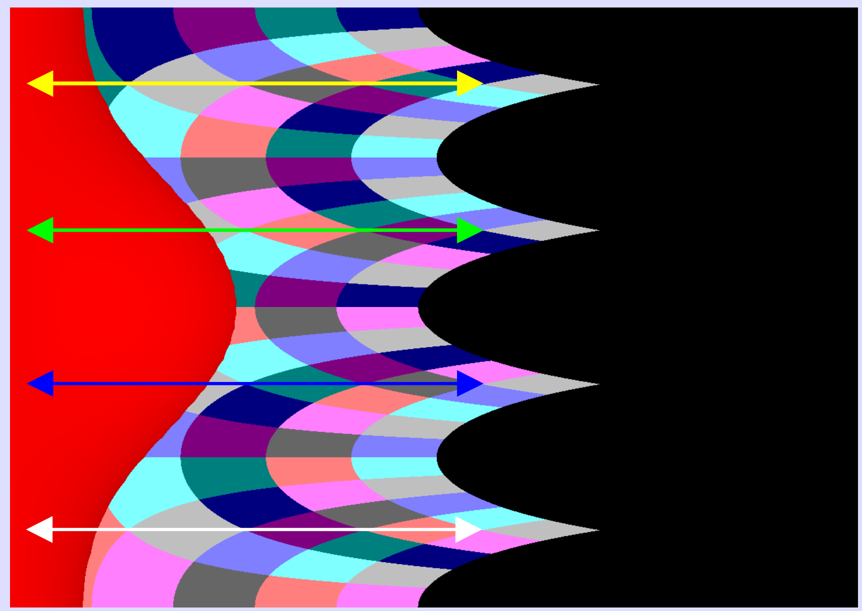

Radial Epipolar Motion Sphere

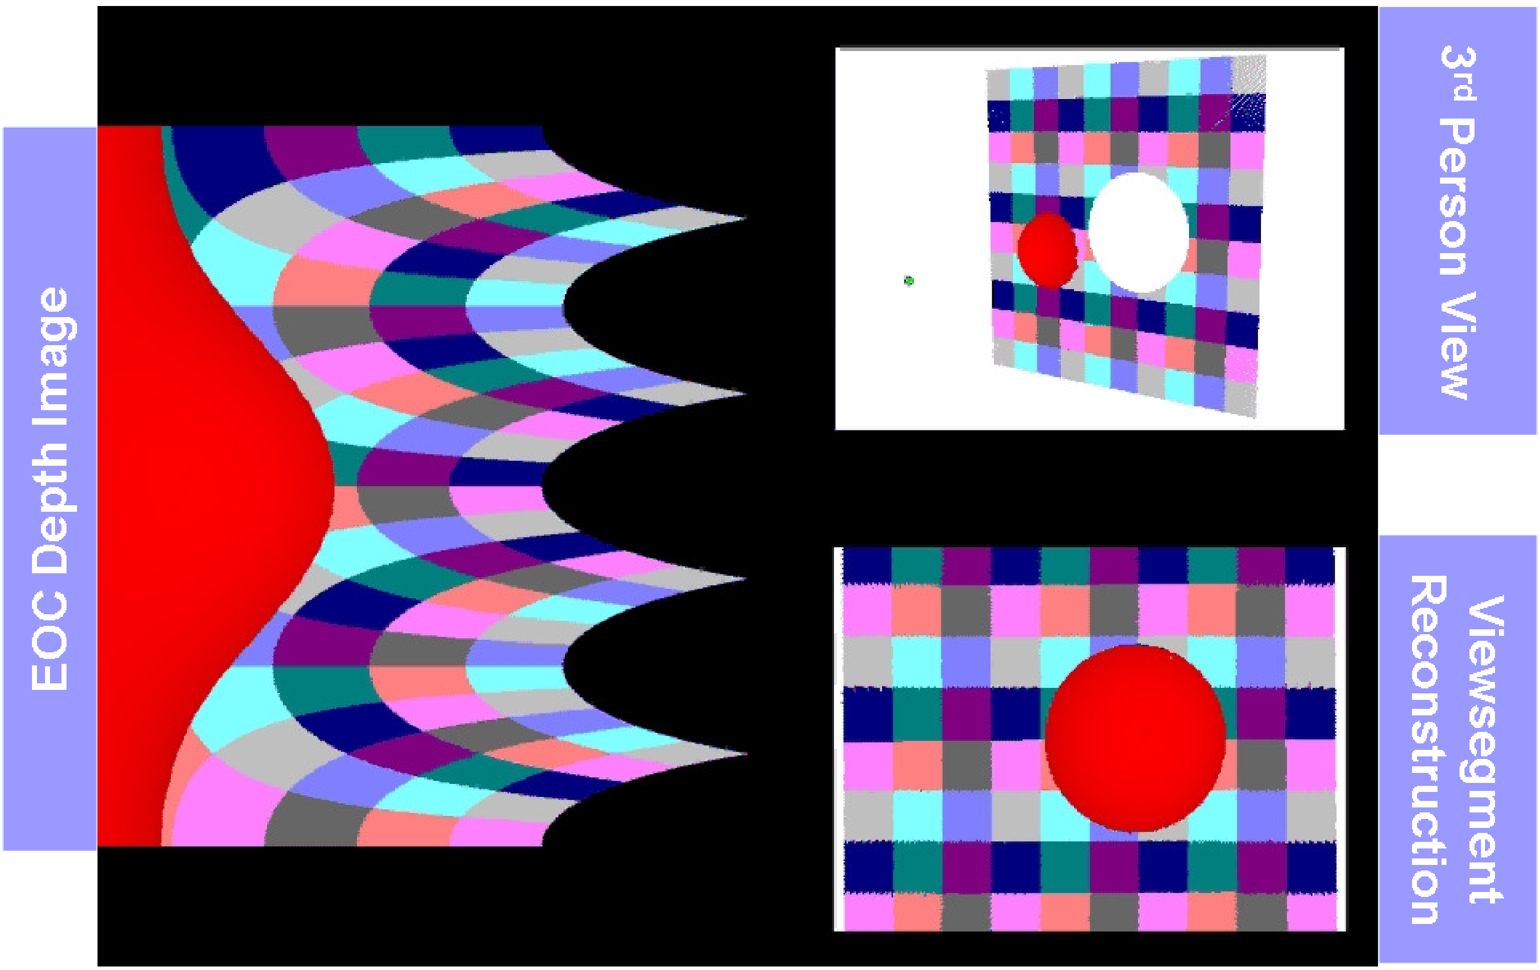

Radial Epipolar Motion
Armadillo

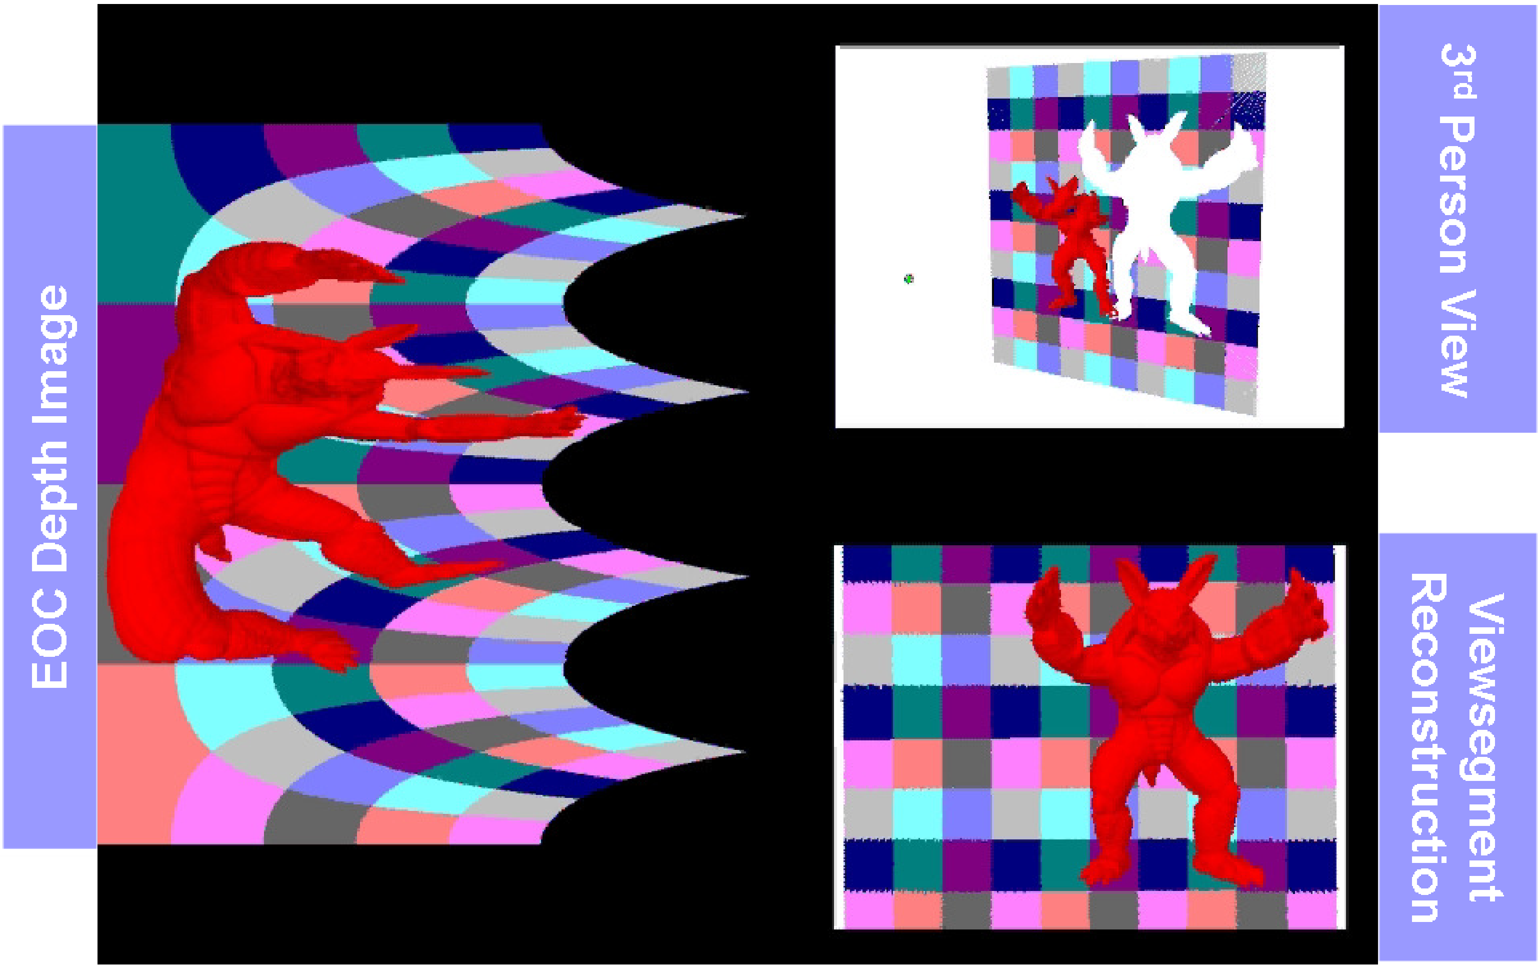

Multiple Images
Crossing Segments of Projection

4 Depth Images	2 EOC Images
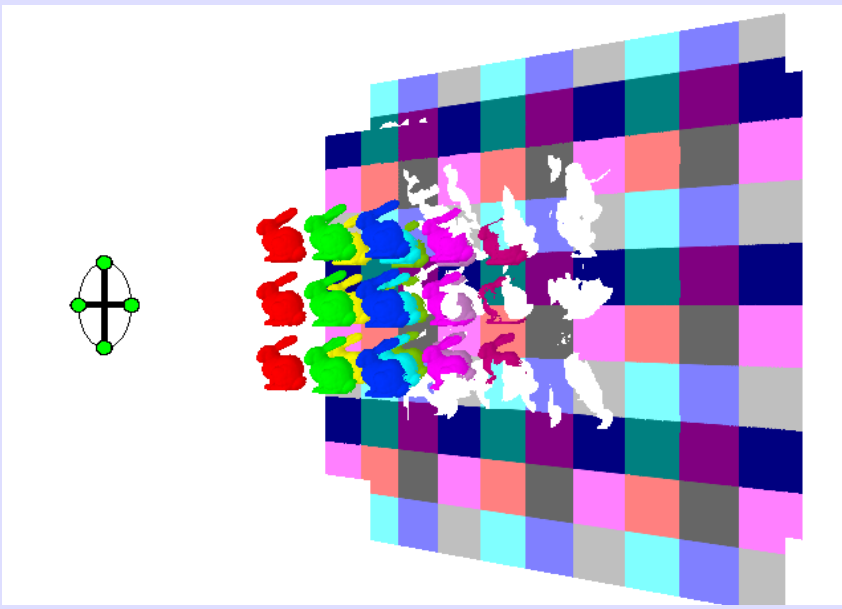	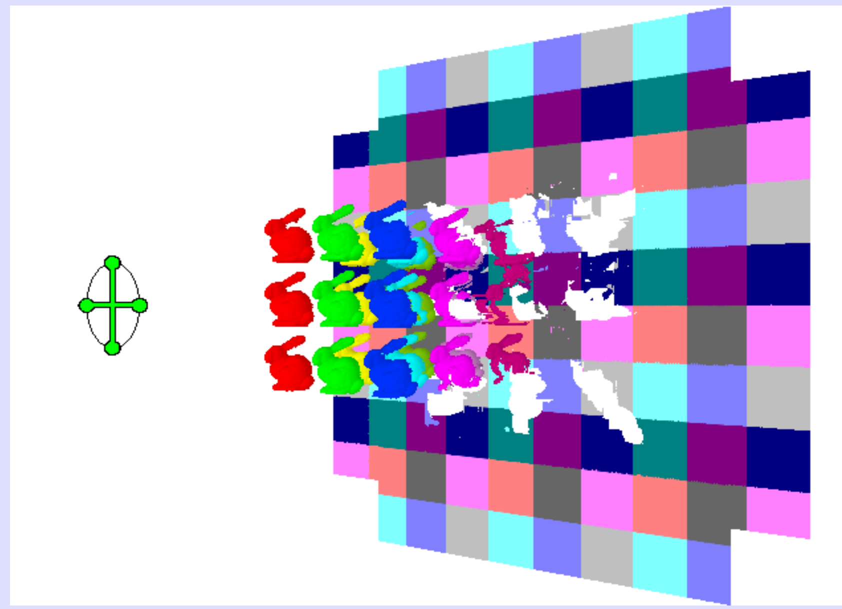

Multiple Images
Crossing Segments of Projection

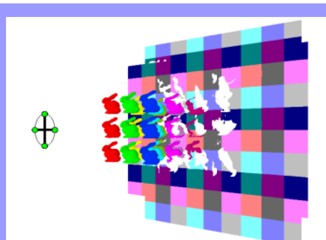

4 Depth Images

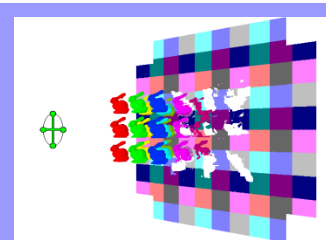

2 EOC Images

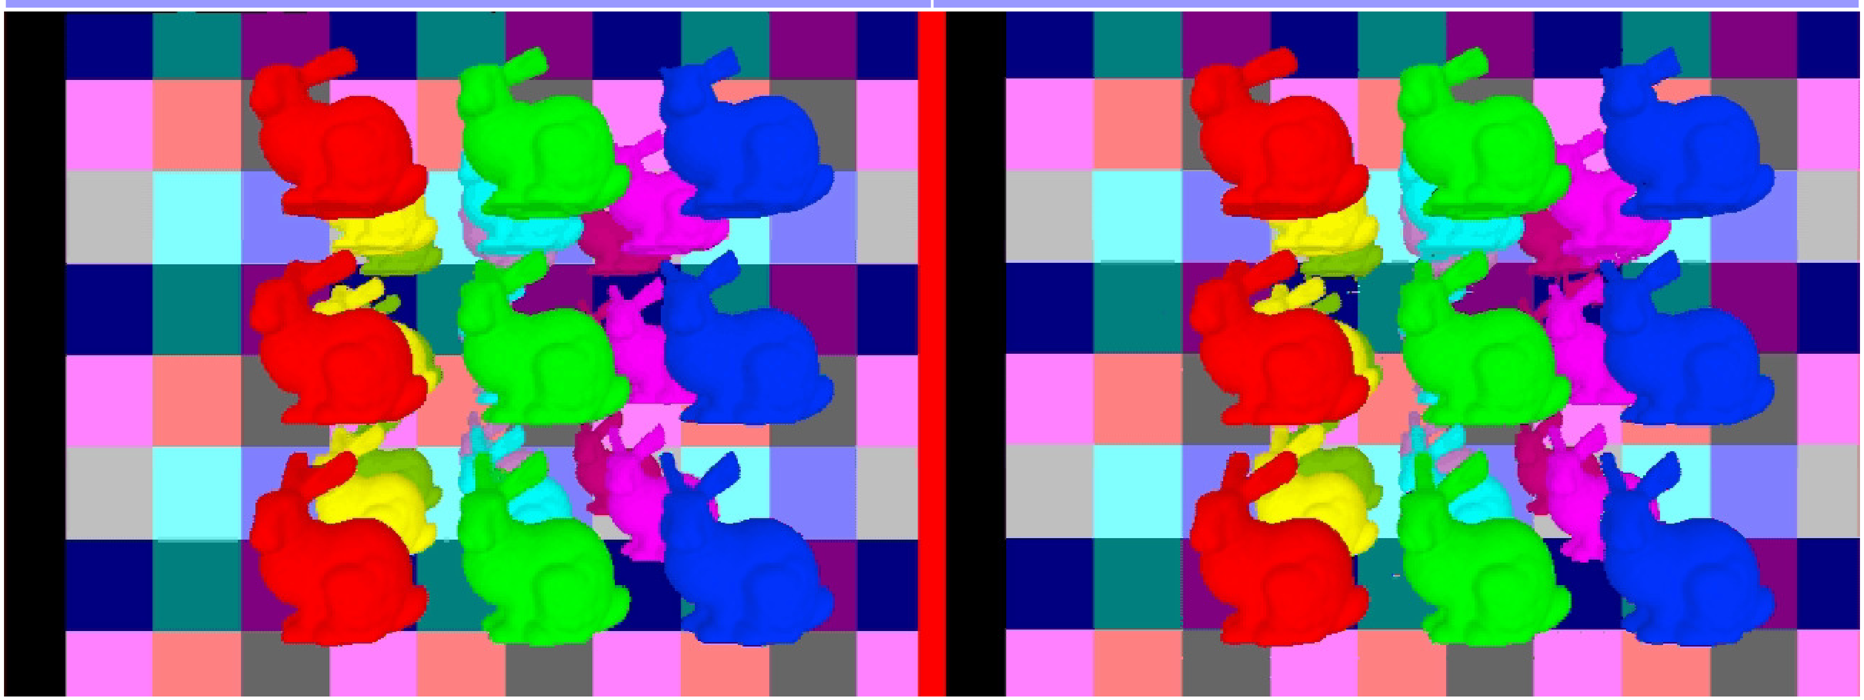

Unity scene

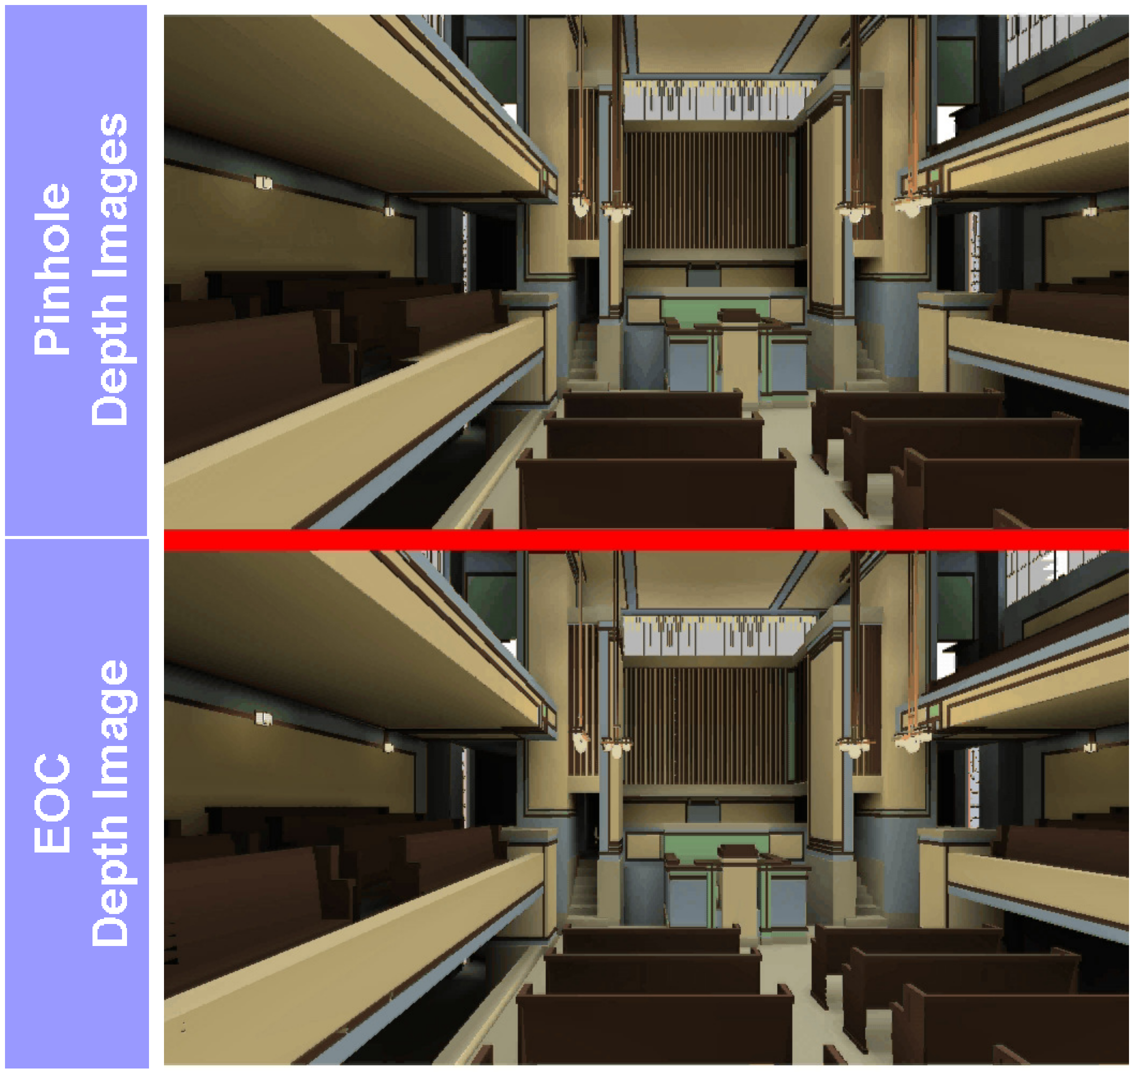

Performance

Scene	Triangles	Construction Time	Rendering Time
Teapot	1K	~100 ms	19 ms
Picket Fence	1K	~100 ms	16 ms
Sphere	1K	~100 ms	36 ms
Armadillo	346K	~100 ms	1,130 ms
Unity	110K	~100 ms	1,793 ms
Bunny	1K 4K 16K 69K	~100 ms	15 ms 20 ms 25 ms 68 ms

Conclusions

- The EOC is a robust solution to disocclusion errors in complex scenes
  - Not conservative, as only the first layer depth discontinuities are examined
- Epipolar constraints are leveraged to separate disocclusion errors in the image plane into a manageable set of 1-D disocclusion errors

Future work

- The EOC is an infrastructure-level innovation
  - Produces a “better” depth image
  - Many potential applications
    - Impostors, simplification, LOD
    - Visibility, occlusion culling
    - 3-D environment maps for reflections & refractions
    - Compression, 3-D display bandwidth reduction
- The EOC generalizes the viewpoint (0-D) to a view*segment* (1-D)
  - Generalizations to a view*triangle* (2-D) or a view*tetrahedron* (3-D)
- Constructing occlusion camera images from photographs?
  - Might be possible without an intermediate 3-D reconstruction

Future work

- Apprehension about removing pinhole constraint not justified
  - Non-pinhole images can be rendered quickly in feed-forward fashion
  - Fast projection is not reserved for pinholes
- Effective and efficient non-pinholes open the door to *camera model design*
  - Abandon conventional simplicity and rigidity of camera model
  - Design and dynamically optimize camera model for application and data set at hand

Acknowledgments

- NSF grant CNS-0417458
- Members of the Computer Graphics and Visualization Lab @ Purdue

Questions