AUGMENTED REALITY USING FULL PANORAMIC CAPTURED SCENE LIGHT-DEPTH MAPS

Aldo René Zang, Luiz Velho
Visgraf Laboratory - IMPA
zang@impa.br, lvelho@impa.br

Dalai Felinto
Fisheries Centre, UBC
d.felinto@fisheries.ubc.ca

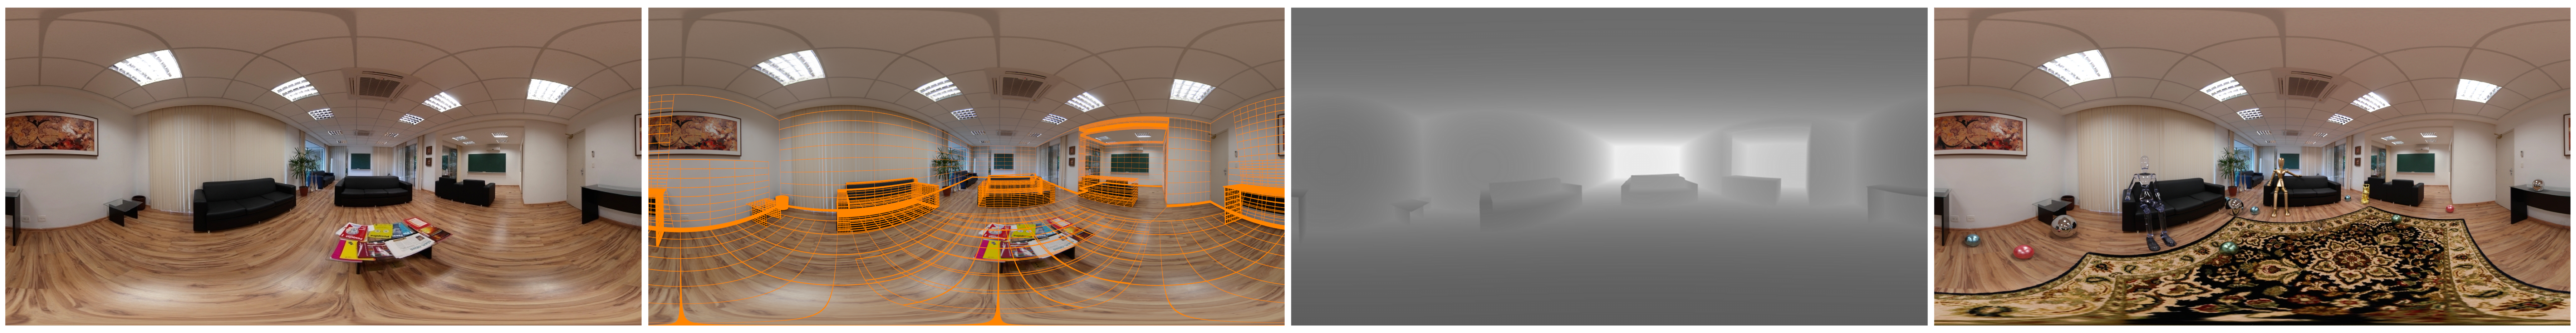

Figure: Left: Captured scene radiance; **Center:** Environment mesh (left) and depth image from modeled environment (right); **Right:** augmented scene rendering.

ABSTRACT

Augmented reality rendering with the reconstructed depth of the scene incorporated in a light-depth map for more accurate shadows and reflections of synthetic elements into a full panoramic scene.

PRODUCTION PIPELINE

1. Environment capture and panorama assembling
2. Calibration of the equirectangular panorama
3. Environment modeling
4. Creation of the light-depth environment map
5. Modeling of the new objects
6. Rendering of the augmented reality panorama

LIGHT-DEPTH ENVIRONMENT MAP

A *light-depth map* contains both radiance and the spatial displacement (i.e., depth) of the environment light.

The *light-depth map* can be constructed from an *HDR* environment map by adding the depth channel. The depth channel can be obtained by a special render from the reconstructed environment meshes, by 3d scanning the environment or other techniques.

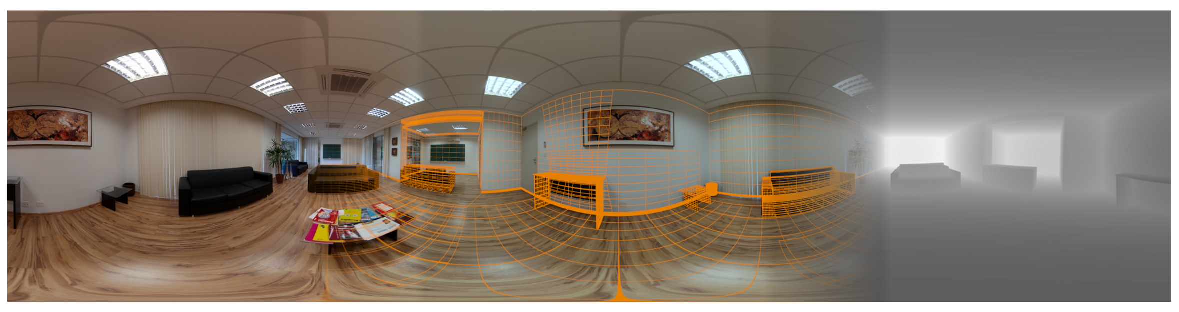

Figure 1: From left to right you can see the environment map, modeled mesh and depth image used to compose the *light-depth map*.

A pixel sample from the *light-depth environment map* is denoted by $M(\omega_i, z_i)$, where ω_i is the direction of the light sample in the map and the scalar value z_i denotes the distance of the light sample from the light space origin.

RENDERING PIPELINE FOR AUGMENTED SCENES

Computing Shadows

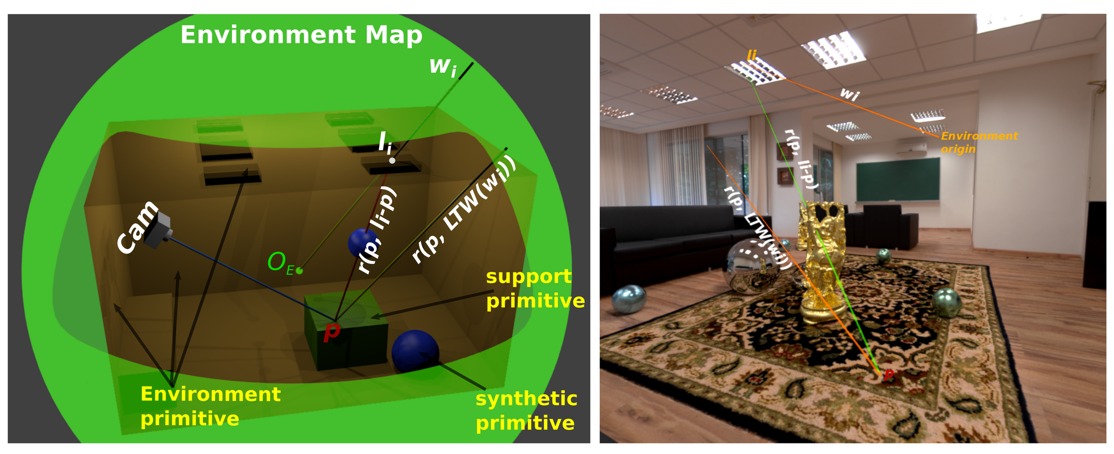

Figure 2: The ray $r(p, LTW(\omega_i))$ used by the directional approach is not occluded, so the light l_i illuminates the point p. The ray $r(p, l_i - p)$ used by our light-positional approach is occluded by a synthetic object and gets the correct world-based estimation.

Computing Reflections

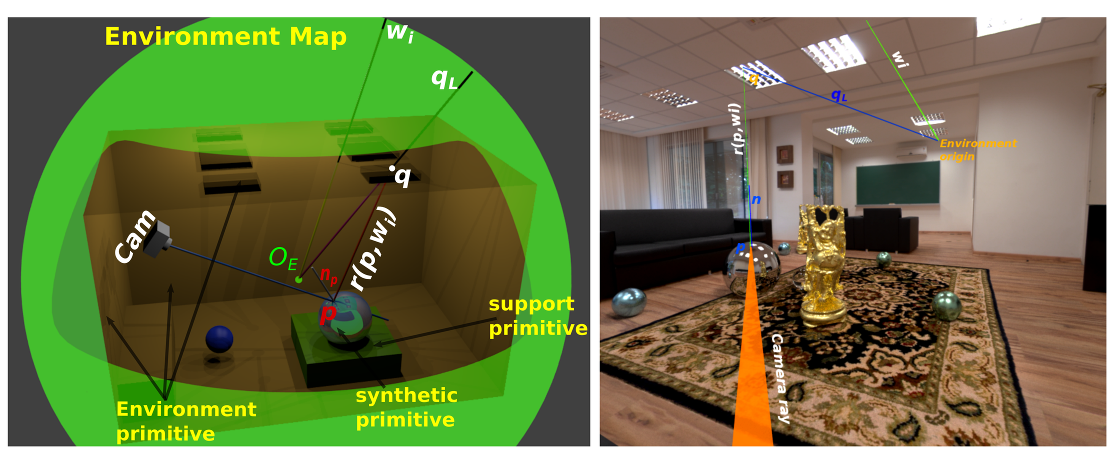

Figure 3: The ray $r(p, \omega_i)$ intersects the scene at point q, over the environment mesh. The radiance of q is stored at q_L direction in the environment map. Our approach computes the correct reflection q_L, instead of computing the ray ω_i that would give a wrong reflection mapping.

Primitives

- ***Synthetic primitives:*** objects that are new to the scene. They don't exist in the original environment.
- ***Support primitives:*** surfaces present in the original environment that needs to receive shadows and reflections from the synthetic primitives.
- ***Environment primitives:*** all the surfaces of the original environment that need to be taken into account for the reflections and shadows computations for the other primitive types.

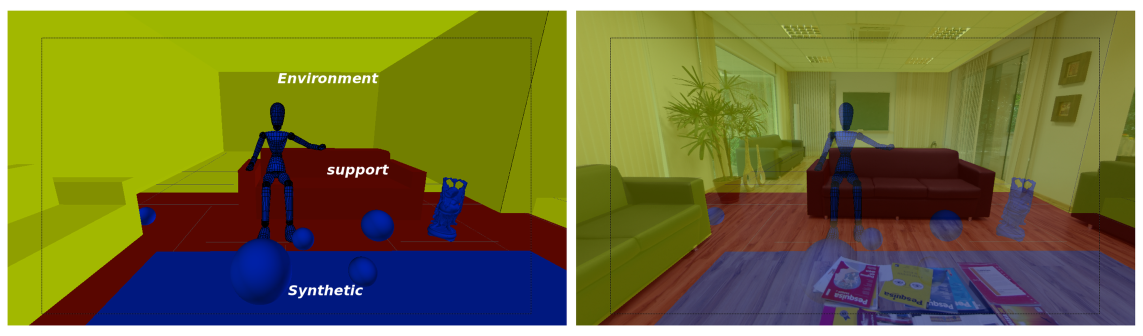

Figure 4: Primitives classification: synthetic (blue), support (red) and environment (olive) primitives.

Surface Scattering: Estimating the Direct Lighting

- ***Synthetic primitives:*** is used the traditional Monte Carlo estimator

$$L_o(p,\omega_o) = \frac{1}{N}\sum_{j=1}^{N} \frac{f(p,\omega_o,\omega_j)L_d(p,\omega_j)|\cos\theta_j|}{pdf(\omega_j)} \tag{1}$$

- ***Environment primitives:*** direct lighting contribution is obtained directly from the *light-depth map* as

$$L_o(p,\omega_o) = M(\omega_p/|\omega_p|, |\omega_p|), \tag{2}$$

where $\omega_p = WTL(p)$ transform p from world space to light space.

- ***Support primitives:*** $L_o(p,\omega_o)$ is computed using the estimator

$$L_o(p,\omega_o) = \frac{1}{N}\sum_{j=1}^{N} \frac{M(\frac{\omega_p}{|\omega_p|}, |\omega_p|)\cdot ES(p,n_p)\cdot \left\langle \frac{\bar{\omega}_j}{|\bar{\omega}_j|}, n_p \right\rangle}{p(\omega_j)}, \tag{3}$$

$ES(p, n_p)$ is a scale factor term that represents the percentage of light contribution coming from the map to the point p.

$$ES(p,n_p) = \frac{\int_{S^2} Lum(M(\omega_s,z_s))\mathrm{d}\omega_s}{\int_{S^2} Lum(M(\omega_j,z_j))\left\langle \frac{\bar{\omega}_j}{|\bar{\omega}_j|}, n_p \right\rangle \mathrm{d}\omega_j}, \tag{4}$$

where $\bar{\omega}_j = LTW(z_j\omega_j) - p$, LTW is the light to world transformation.

AUGMENTED REALITY PATH TRACING

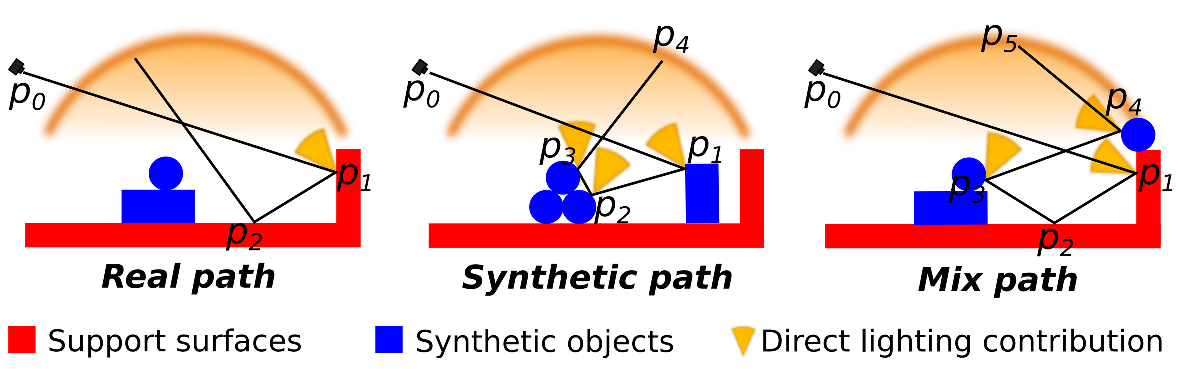

Figure 5: The yellow cones shows which vertices contributes for direct lighting (DL) account. In real paths only the first vertex contributes with DL. Synthetic path are computed normally. Mixed paths contributes only at first vertex p_1, at all synthetic vertices, and for all support vertices that are the next path step of a synthetic vertex.

There is no need to calculate any of the real path for the light. They are already present in the original photography. However, since part of the local scene was modeled (and potentially modified), we need to re-render it for these elements. The calculated radiance needs to match the captured value from the environment. We do this by aggregating all the real path radiance in a vertex p_1 as direct illumination, discarding the need of considering the neighbor vertices light contribution.

The synthetic and mixed light paths are calculated in a similar way, taking the individual path vertex light contributions for every vertex of the light path. The difference between them is in the Monte Carlo estimate applied in each case. In the figure (5) you can see the different light path types and the calculation that happens on the corresponding path vertices.

RESULTS

Figure 6: Comparison between reflections given by directional (left) and light-depth (right) maps approaches.

Figure 7: Shadows computed using the *light-depth environment map*. The orientation of the shadows varies according to the objects position.

Figure 8: Shadows and reflections computed using the *light-depth environment map*. It's possible to see how the highlights and shadows vary with the position of the balls in the scene.

Figure 9: Camera traveling effect. Two points of view using a camera position different than the one used to capture the environment. A wall partially visible in left image is occluded in the right image.

Figure 10: The couch is real and has been stored on the light-depth map. We deformed the environment texture coordinates to simulate the interaction between the synthetic spheres and the real couch.

References

1. Debevec, P. 1998. Rendering synthetic objects into real scenes: bridging traditional and image-based graphics with global illumination and high dynamic range photography. In SIGGRAPH '98, ACM, New York, NY, USA, SIGGRAPH '98, 189–198.
2. Felinto, D. Q., Zang, A. R., and Velho, L. (2012). Production framework for full panoramic scenes with photo-realistic augmented reality. In XXXVIII Latin American Conference of Informatics (CLEI), Medellin.
3. Zang, A. R. and Velho, L. (2011a). Arluxrender project. http://www.impa.br/~zang/arlux. Visgraf.