

# Laboratório VISGRAF

Instituto de Matemática Pura e Aplicada

## **Global Illumination of Non-Euclidean Spaces**

*Tiago Novello, Vinicius da Silva, Luiz Velho*

Technical Report    TR-20-06    Relatório Técnico

April - 2020 - Abril

The contents of this report are the sole responsibility of the authors.  
O conteúdo do presente relatório é de única responsabilidade dos autores.

# GLOBAL ILLUMINATION OF NON-EUCLIDEAN SPACES

TIAGO NOVELLO, VINICIUS DA SILVA, AND LUIZ VELHO

ABSTRACT. This paper presents a path tracer algorithm to compute the global illumination of non-Euclidean manifolds. We use the 3D torus as an example.



## 1. INTRODUCTION

*Global illumination* is a collection of algorithms in computer graphics that are employed to mimic realist lighting in 3D scenes. The complexity of these algorithms depends on the *indirect illumination*, which comes not directly from the light sources (*direct illumination*) but the light of these sources reflected by other surfaces at the scene. Then computing *reflections*, *refractions*, and *shadows* are important to compute the global illumination of a given scene.

In *photorealistic rendering* a set of techniques are used to create an image from a 3D scene that approximates from a photograph [4], thus enabling the visualization of the global illumination of the underlying 3D scene. *Path tracing* is an algorithm contained in this set. Informally, it consists of giving color for each point (*eye*) and direction (*pixel*) by launching a path of rays through the scene based on reflections, refractions, and shadows. To write a ray tracer it is common to simulate *cameras*, *ray-surface intersections*, *light sources*, *visibility*, *surface material*, *indirect light transport*, and *ray propagation*.

The above concepts used in global illumination of Euclidean spaces can be extended to non-Euclidean spaces. These abstract spaces were developed avoiding Euclid's fifth postulate, they do "not" exist in the real world. Therefore no real

---

*Date:* June 24, 2019.

*Key words and phrases.* Non-Euclidean geometry, path tracer.

camera can take a picture in such spaces, which makes exploring photorealistic rendering very challenging. Attempts to visualize non-Euclidean geometries have produced beautiful and inspiring images with no counterpart in the real world; much of this is due to the non-trivial topology and geometry of such abstract spaces.

Berger et al. [1] was the first to proposed an *image-based* algorithm to visualize non-Euclidean spaces. Their rendering algorithm exploited programmable compute shaders and CUDA to implement ray-tracing on the GPU. However, the algorithm was limited to render only scenes composed of Lambertian surfaces described as implicit surfaces. Shadows and reflections were not simulated. In [6], we proposed the first real-time visualization of such spaces using RTX GPUs. More complex scenes were rendered, and shadows and reflections were introduced (Figure 1).

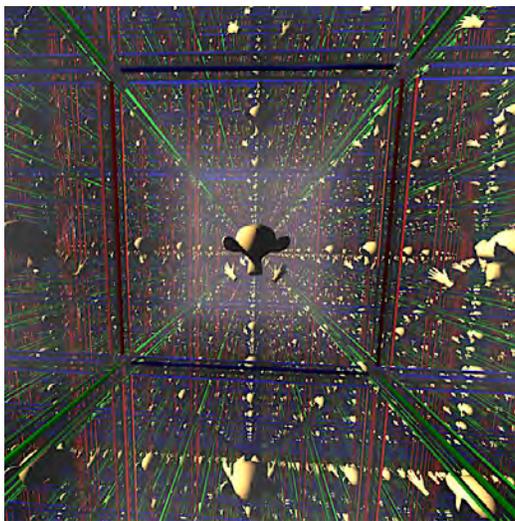


FIGURE 1. Inside view of the 3D flat torus using ray tracing [6].

We take the challenge of computing the global illumination in non-Euclidean spaces using a path tracer. In such spaces, the concept of photorealistic rendering is not well defined because there is no real camera there. However, the possibility of defining a *global illumination equation* in non-Euclidean geometry opens the door for a path tracing generalization. The images are inspiring.

### 1.1. Non-Euclidean Spaces.

In dealing with path tracing algorithms is required: (1) Space must be locally similar to a Euclidean space — a *manifold*. This allows us to model the camera and scene objects; (2) For each point  $p$  we need *tangent* vectors pointing in all directions. The *inner product* between two tangent vector is required to simulate light effects; (3) For a point  $p$  and a vector  $v$  tangent at  $p$ , we need a *ray*, and its intersections with objects in the scene.

*Geometric manifolds* satisfies the above properties. Such objects are locally geometrical similar to special spaces called *model geometries*. In dimension two, for example, there are exactly three models: Euclidean, hyperbolic, and spherical spaces. In dimension three, there are five more model geometries, however, in this work, we focus on the (classical) first three spaces. We describe these topics in more detail below. Great texts on this subject are Thurston [5] and Martelli [3].

**Definition 1.1** (Euclidean space). The *Euclidean space*  $\mathbb{E}^3$  is  $\mathbb{R}^3$  endowed with the classical *inner product*  $\langle u, v \rangle_{\mathbb{E}} = u_x v_x + u_y v_y + u_z v_z$  where  $u$  and  $v$  are vectors in  $\mathbb{R}^3$ . The *distance* between two points  $p$  and  $q$  is  $d_{\mathbb{E}}(p, q) = \sqrt{\langle p - q, p - q \rangle_{\mathbb{E}}}$ . The curve  $\gamma(t) = p + tv$  describes a *ray* leaving a point  $p$  in a direction  $v$ . Analogously, for any  $n > 0$  the Euclidean space  $\mathbb{E}^n$  is constructed.

**Definition 1.2** (Hyperbolic space). The *Lorentzian inner product* of the vectors  $v$  and  $u$  in  $\mathbb{R}^4$  is  $\langle u, v \rangle_{\mathbb{H}} = u_x v_x + u_y v_y + u_z v_z - u_w v_w$ . The *Lorentzian space* is  $\mathbb{R}^4$  endowed with this product. The *hyperbolic space*  $\mathbb{H}^3$  is the hyperboloid  $\{p \in \mathbb{R}^4 \mid \langle p, p \rangle_{\mathbb{H}} = -1\}$  endowed with a special metric  $d_{\mathbb{H}}(p, q) = \cosh^{-1}(-\langle p, q \rangle_{\mathbb{H}})$ .

A tangent vector  $v$  to a point  $p$  in  $\mathbb{H}^3$  satisfies  $\langle p, v \rangle_{\mathbb{H}} = 0$ , then the *tangent space*  $T_p \mathbb{H}^3 = \{v \in \mathbb{R}^4 \mid \langle p, v \rangle_{\mathbb{H}} = 0\}$ . *Rays* in  $\mathbb{H}^3$  are the intersections between  $\mathbb{H}^3$  and the planes in  $\mathbb{R}^4$  containing the origin. For instance, the ray leaving a point  $p \in \mathbb{H}^3$  in a tangent direction  $v$  is the intersection between  $\mathbb{H}^3$  and the plane spanned by the vectors  $v$  and  $p$  in  $\mathbb{E}^4$ . Such ray can be parameterized as  $r(t) = \cosh(t)p + \sinh(t)v$ .

$\mathbb{H}^3$  contains no straight line, thus its rays can not be straight. However, it is possible to model  $\mathbb{H}^3$  in the unit open ball in  $\mathbb{R}^3$  — known as *Klein model*  $\mathbb{K}^3$ — such that the rays are straight lines. More precisely, each point  $p \in \mathbb{H}^3$  is projected in the space  $\{(x, y, z, w) \in \mathbb{R}^4 \mid w = 1\}$  by considering  $p/p_w$ .

**Definition 1.3** (Elliptic Space). The *3-sphere*  $\mathbb{S}^3$  is the set  $\{p \in \mathbb{E}^4 \mid \langle p, p \rangle_{\mathbb{E}} = 1\}$  endowed with the metric  $d_{\mathbb{S}}(p, q) = \cos^{-1} \langle p, q \rangle_{\mathbb{E}}$ . A tangent vector  $v$  to a point in  $\mathbb{S}^3$  satisfies  $\langle p, v \rangle_{\mathbb{E}} = 0$ , then the *tangent space*  $T_p \mathbb{S}^3 = \{v \in \mathbb{E}^4 \mid \langle p, v \rangle_{\mathbb{E}} = 0\}$ . The space  $T_p \mathbb{S}^3$  inherits the Euclidean inner product of  $\mathbb{E}^4$ . A *ray* in  $\mathbb{S}^3$  passing through a point  $p$  in a tangent direction  $v$  is the intersection between  $\mathbb{S}^3$  and the plane spanned by  $v$ ,  $p$ , and the origin. Such ray can be parameterized as  $r(t) = \cos(t)p + \sin(t)v$ .

A *3-manifold*  $M$  is a topological space which is locally identical to the Euclidean space  $\mathbb{E}^3$ : there is a neighborhood of every point in  $M$  mapped diffeomorphically to the open ball of  $\mathbb{E}^n$ . These maps are called *charts*. The change of charts between two neighborhoods in  $M$  must be differentiable. Examples of 3-manifolds include the Euclidean, hyperbolic, and spherical spaces.

*Metric* and *rays* are fundamental objects when working with path tracing, since light travels along with rays and are distributed in the scene surfaces by the metric. A manifold  $M$  endowed with a metric at each tangent space (in a smooth manner) is called *Riemannian*. This allows us to compute the length of *vectors* and distances between points in  $M$ . A *geodesic* (ray) in  $M$  is a curve locally minimizing length.

We remember an algebraic (very computational) way to construct manifolds from simple ones. Let  $M$  be a connected manifold and  $\Gamma$  be a discrete group of isometries acting on  $M$ , the *quotient*  $M/\Gamma$  is the set  $\{\Gamma \cdot p \mid p \in M\}$  where  $\Gamma \cdot p = \{g(p) \mid g \in \Gamma\}$  is the *orbit* of  $p$ . For example, the quotient of  $\mathbb{E}^3$  by the group of translation gives rise to the *flat torus*  $\mathbb{T}^3$ . A ray  $r$  leaving a point  $p \in \mathbb{T}^2$  in a direction  $v$  is described by considering the fractional part of the coordinates of  $r(t) = p + t \cdot v$ .

Interesting cases arise when  $M$  is a *geometry* model: Euclidean, hyperbolic, and spherical. The space  $M/\Gamma$  inherits the *geometric structure* of  $M$ . For example,  $\mathbb{T}^3$  has the geometric structure modeled by  $\mathbb{E}^3$ . The *fundamental domain*  $\Delta$  plays an important role in the above construction being the region of  $M$  containing exactly one point for each orbit. The unit cube is the fundamental domain of  $\mathbb{T}^3$ .

## 2. ILLUMINATION OF NON-EUCLIDEAN SPACES

This section focus on extending the local and global illumination equations to a Riemannian manifold  $(M, g)$ , the next section presents a path tracer algorithm for the class of geometric manifolds. We consider that light propagates along with rays in  $(M, g)$ , and that it is constant in a vacuum. Then inside views of a scene inside  $M$  can be rendered by tracing rays: given a point (eye) and a direction (pixel), we trace a geodesic (first ray). When it reaches a point on a surface object, we compute its *Riemannian illumination* by considering the *direct* and *indirect* light contributions. The direct illumination considers the rays coming directly from light sources and the indirect gathers the light rays bounced by other surfaces.

*Illumination* is the process of simulating the light leaving a point in a given direction. Classic approaches to perform such tasks are not suited for Riemannian manifolds, so we propose a more general illumination definition.

Let  $\mathcal{C}$  be the space of RGB colors. The *Riemannian illumination* of a 3-manifold  $M$ , with a scene embedded on it, is a function  $L : M \times \mathbb{S}^2 \rightarrow \mathcal{C}$ , where for each point  $p \in M$  and unit tangent direction  $v \in T_p M$  we have a color  $L(p, v)$  representing the light leaving this direction. Because we are considering the light constant along geodesics, the function  $L$  only needs to be computed on the scene surface objects.

There are many approaches to compute the illumination function. We start with a local model considering only the direct contributions of the light sources and surfaces composed of diffuse and specular materials. Then we incorporate the indirect light contribution generating a global model.

## 2.1. Riemannian local illumination.

We start with the classical local illumination in a Euclidean scene. Let  $p$  be a point in a surface object  $S$ , and  $N$  its normal vector. Let  $l$  be a point light,  $w_i = (l - p)/|l - p|_{\mathbb{E}}$  is the *incident vector* and  $w_r = -w_i + 2\langle w_i, N \rangle_{\mathbb{E}} N$  is the *perfect reflection* of  $w_i$ . For a vector  $v$  leaving  $p$ , the *local illumination function* is:

$$(2.1) \quad L(p, v) = k_a L_a + k_d \sum_{j=1}^n L_i \langle w_i^j, N \rangle_{\mathbb{E}} + k_s \sum_{j=1}^n L_i \langle w_r^j, v \rangle_{\mathbb{E}}^{k_e},$$

where  $L_a$  indicates the ambient light intensity,  $L_i$  is the radiance emitted by the light source  $l_i$ , and the numbers  $k_a$ ,  $k_d$ ,  $k_s$ , and  $k_e$  represent the *ambient*, *diffuse*, *specular*, and *roughness* coefficients of the surface  $S$  at  $p$ .

We now extend Equation 2.1 to a Riemannian manifold  $(M, g)$ ;  $g$  is the Riemannian metric. Here  $p$  is a point in a smooth surface  $S$  embedded in  $M$ , and  $N$  is the normal vector at  $p$ . The direction  $w_i$  is the unit tangent vector at time zero of the geodesic running from  $p$  towards the light source  $l$ . Then  $w_r = -w_i + 2g_p(w_i, N)N$  is the perfect reflection of  $w_i$ . The metric  $\langle \cdot, \cdot \rangle_{\mathbb{E}}$  used in Equation 2.1 is replaced by the Riemannian metric  $g_p$  of  $M$  restricted to  $p$ , this gives rise to the *Riemannian local illumination* for a unit direction  $v$  at  $T_p M$ :

$$(2.2) \quad L(p, v) = k_a L_a + k_d \sum_{j=1}^n L_i g_p(w_i^j, N) + k_s \sum_{j=1}^n L_i g_p(w_r^j, v)^{k_e},$$

To render an image of a scene embedded in the Riemannian manifold  $M$  we define a *Riemannian ray casting*: it shoots rays through the pixels, if a ray intersects a surface, we compute its color using Riemannian local illumination function. Consider a 2-sphere  $\mathbb{S}_o^2$  centered in a point  $o$  (the observer) in  $M$ . We give a color

for each sphere point (ray direction) in the observer field of view  $V$  by tracing a ray; In other words,  $V \subset \mathbb{S}_o^2$  carries the image. We call this procedure *Riemannian ray casting*. Specifically, the unit sphere  $\mathbb{S}_o^2$  is centered at the origin of  $T_oM$ . For each direction  $v$  in  $\mathbb{S}_o^2$  we attribute a color  $c$  by launching a ray  $\gamma$  from  $o$  in the direction  $v$ . Each time  $\gamma$  intersects a scene object at a point  $p = \gamma(t)$  we compute its Riemannian local illumination  $L(p, -\gamma'(t))$ .

Before we go to a global illumination function we observe in Equation 2.2 that the ambient term  $k_a L_a$  is a very rough approximation of the indirect light contribution. The other terms in Equation 2.2 is the direct light contribution when dealing with diffuse/specular surfaces.

## 2.2. Global illumination in non-Euclidean manifolds.

We remember the classical *global illumination function* of the Euclidean space  $\mathbb{E}^3$  introduced by Kajiyi [2]. For each point  $p$ , it computes the amount of light emitted in a direction  $v$ . It is modeled through the integral equation over the hemisphere  $\Omega(p) = \{v \in \mathbb{S}^2 \mid \langle v, N(p) \rangle \geq 0\}_{\mathbb{E}}$ :

$$(2.3) \quad L(p, v) = L_e(p, v) + \int_{\Omega} f_r(p, v, w_i) L(p, w_i) \langle w_i, N \rangle_{\mathbb{E}} dw_i.$$

The *bidirectional reflectance distribution function* (BRDF)  $f_r(p, v, w_i)$  defines how the light reflects at  $p$ , and  $L_e(p, w)$  is the light emitted at  $p$  in the direction  $v$ .

Equation 2.3 generalizes Equation 2.1 because it integrates over all direction entering the hemisphere  $\Omega$  while Equation 2.1 considers only the directions coming directly from the light points. The other indirect directions are roughly approximated by the ambient term  $k_a L_a$ .

Again, replacing the Euclidean metric with the Riemannian metric, we obtain a global illumination function for a Riemannian manifold  $(M, g)$ ,

$$(2.4) \quad L(p, v) = L_e(p, v) + \int_{\Omega} f_r(p, v, w_i) L(p, w_i) g_p(w_i, N) dw_i.$$

Taking into account the import contribution of the directions  $w_i$  leaving the point  $p$  in the direction of the light points  $l_i$ , we divide Equation 2.4 into the direct and indirect components  $L(p, v) = L_e(p, v) + L_{dir}(p, v) + L_{ind}(p, v)$ . The direct contribution  $L_{dir}(p, v)$  is computed using Riemannian local illumination function (Equation 2.2). For the indirect contribution  $L_{ind}(p, v)$ , we use *Monte Carlo integration* to estimate a good approximation.

Let  $\{w_i^1, \dots, w_i^n\}$  be  $n$  vectors on the hemisphere  $\Omega$  chosen using a distribution density  $d_w$ , the Monte Carlo integration states that

$$(2.5) \quad L_{ind}(p, v) \approx \frac{1}{n} \sum_{k=1}^n \frac{f_r(p, v, w_i^k) L(p, w_i^k) g_p(w_i^k, N)}{d_w(w_i^k)}.$$

If  $L(p, w_i^k)$  is known, the *law of large number* ensures the convergence when  $n \rightarrow \infty$ . Otherwise, computing these approximation for each  $L(p, w_i^k)$ , and repeating such procedure for a finite number  $m$  of times, we obtain a version of the *path tracing* algorithm. It is well-known in compute graphics, since Kajiyi [2], that the number  $m$  does not need to be large for computing realistic images.

### 3. PATH TRACING IN NON-EUCLIDEAN SPACES

We combine Riemannian ray casting and Riemannian global illumination to synthesize inside views of non-Euclidean spaces. This generalizes the classical *path tracing*. We focus on a geometric manifold  $M/\Gamma$  since it admits a well-behaved geometry  $M$  and a combinatorial description of its topology in terms of  $\Gamma$ .

Our method approximates the Riemannian global illumination of the visible surfaces inside a geometric manifold using the ray tracing capabilities of the RTX platform. We will discuss first the basic principles of path tracing in Non-Euclidean spaces, as well as, the general algorithm in CPU. We will show how to map the computation to the RTX pipeline and present the details of GPU implementation.

#### 3.1. Overview of the Method.

The path tracing is the most natural method to produce photorealistic images of Euclidean scenes. Thus extending the technique to non-Euclidean spaces would provide a flavor of the photo-realism inside such abstract spaces. It is necessary to take into account the geometry/topology of the non-Euclidean space. The first aspect of this task is to simulate the ray path as it travels inside the space, starting from the point of observation until it intersects with a visible object — the Riemannian ray casting. The second aspect amounts to the illumination, evaluating the light scattered from the environment in the ray direction — the Riemannian global illumination. Because of the non-trivial topology, the ray path is updated as it exits the fundamental domain.

#### 3.2. Algorithm in CPU.

We present the basic path tracing algorithm for a geometric manifold  $M/\Gamma$ , and compare it with the traditional path tracing of Euclidean space.

We start with an algorithm to trace rays inside  $M/\Gamma$ . Let  $p$  be a point inside the fundamental domain  $\Delta$  of  $M/\Gamma$ , and  $v$  be a direction at  $T_pM$ , Algorithm 1 trace a ray  $\gamma$  from  $p$  towards the direction  $v$ . If  $\gamma$  intersects a scene object in  $\Delta$ , the point and direction are updated.

---

#### Algorithm 1: Trace ray

---

**Data:** point  $p$ , direction  $v$

**Result:** bool  $hit$

Trace a ray  $\gamma$  from  $(p, v)$  inside  $\Delta$ ;

**repeat**

    Find closest intersection  $\gamma(t)$  with objects  $O$  in  $\Delta$ ;

**if**  $\gamma(t) \neq \emptyset$  **then**

        Update  $p = \gamma(t)$  and  $v = \gamma'(t)$ ;

**return** true;

**else**

        Find intersection of  $\gamma$  with faces of  $\Delta$  ;

        Compute the new origin  $p'$  and ray  $\gamma'$ ;

        ++  $i$ ;

**end**

**until**  $i \leq maxlevel$ ;

**return** false;

---

As it can be verified in Algorithm 1, the rays are intersected with visible objects (line 5) and if there is a hit (line 6), the ray point and direction is updated (line 7). The whole computation has the *fundamental domain* as a base, which is modeled by a polyhedron  $\Delta$ . Therefore, as the ray hits a face  $F$  of  $\Delta$  (line 9), we transport it using the corresponding transformation of the discrete group  $\Gamma$  (line 10–11). This is the most important and critical step since it depends on the geometry and topology of the space. As such, it is specific for each type of space. For practical computational reasons, we cannot continue the ray indefinitely, thus a maximum level is set to stop the path (line 15).

Algorithm 2 describes the path tracer. The rays are generated from the observer’s point of view (lines 1–2). Using Algorithm 1 the visible points are computed and if there is a hit (line 3), shading is done using direct and indirect illumination (line 4–5). Note that Algorithm 2 is computing the Riemannian ray casting. To compute a good approximation of the image, it is common to accumulate the resulting image of Algorithm 2 and line 2 chooses the direction associated with the pixel using a distribution density function.

---

**Algorithm 2:** Path Tracing in non-Euclidean spaces

---

```

for each pixel  $\sigma \in I$  do
  Let  $p := 0$  and  $v$  be the direction associated to  $\sigma$ ;
  if  $\text{TraceRay}(p, v, d)$  then
    Define a depth  $d$  and a color  $c = 0$ ;
    Shade  $\sigma$  using  $L_{dir}(p, v) + L_{ind}(p, v, d, c)$ ;
  end
end

```

---

The direct illumination is computed using Equation 2.2. To add the shadows we multiply the direct illumination by a factor 0 if tracing a ray from the hit point to the light source a scene intersection occurs, and 1 otherwise.

The indirect illumination is presented in Algorithm 3. An integer variable *depth*  $d$  is considered to control the number of ray light bounces of the path tracer. If  $d$  is greater than zero and the ray launched from  $(p, v)$  intersect the scene (line 3), we sample a hemisphere direction and compute an approximation of the Riemannian global illumination (line 4–5). Clearly, in line 5 Algorithm 3, is called again. This came from Equation 2.5 considering only one sample.

---

**Algorithm 3:** Indirect illumination

---

```

Data: point  $p$ , direction  $v$ , depth  $d$ , color  $c$ 
Result: color  $c$ 
if  $d \geq 0$  and  $\text{TraceRay}(p, v)$  then
  Sample a direction  $w \in \Omega(p)$  using a distribution  $d_w$ ;
   $c += L_{dir}(p, v) + \frac{f_r(p, v, w)L_{ind}(p, w, d - 1, c)g_p(w, N)}{d_w(w)}$ ;
end
return  $c$ ;

```

---

### 3.3. RTX Pipeline.

NVidia RTX is a hardware/software platform with support for real time ray tracing. The ray tracing code of an application using this architecture consists of CPU host code, GPU device code, and the memory to transfer data to the Acceleration Structures for fast geometry culling when intersecting rays with scene objects. Specifically, the CPU host code manages the memory flow between devices, sets up, controls and spawn GPU shaders and defines the Acceleration Structures.

These shaders correspond to tasks in Algorithm 1, 2, and 3. The *Ray Generation Shader* is responsible for creating the rays (line 1 in Algorithm 2), which are defined by their origins and directions (line 2). A call to `TraceRay()` launches a ray (line 3 in Algorithm 1). The next stage is a fixed traversal of the Acceleration Structure described only at a high level here. This traversal uses an *Intersection Shader* to calculate the intersections (line 5 in Algorithm 1). All hits found pass by tests to verify if they are the closest hit. After no additional hits are found, the *Closest-Hit Shader* is called for the closest intersection point (line 7 in Algorithm 1). In case no hits are found, Miss Shader is called as a fallback case. It is important to note that additional rays can be launched in the Closest-Hit and Miss shaders.

The *scene objects* and the *boundary of the fundamental domain* are treated differently when mapping the algorithm to the RTX pipeline — while the scene objects are tested and shaded in the regular way (lines 5 and 7 in Algorithm 1), the boundary of the fundamental domain is used to transport the rays by the discrete group (lines 9 and 10). This is implemented with a custom designed Miss Shader.

### 3.4. GPU Implementation.

The implementation of our visualization platform in GPU is built on top of Falcor using DirectX 12 on Windows 10. The Falcor development framework consists of a library with support for DXR at a high level and a built-in scene description format. We use the software Blender to create the scene objects.

The core functionality of our system’s architecture consists of a set of shaders that are mapped to the RTX GPU pipeline as described above. We developed generic shaders for each stage of the GPU path tracing pipeline that are independent of the geometric structure of the non-Euclidean space.

**Ray Generation Shader:** Creates camera rays using the isometries of the space to transform the ray origin and direction to the camera coordinate system.

**Intersection Shader:** Computes the ray-object intersection using a parameterization of the ray. Ray and objects are defined based on the model Geometry.

**Closest Hit Shader:** Performs the shading operation. This includes computing the local and global illumination. The local illumination amounts to the direct contribution of light sources that are based on angles between the light direction and the surface normal, as well as, the distance to the light. The global illumination is computed by launching indirect rays based on the surface BRDF.

**Miss Shader:** Deals with the transport of rays in the covering space, as they leave the fundamental domain. For this, the rays are tested for intersection with the boundary of the polyhedron  $\Delta$ .

## 4. EXAMPLES AND RESULTS

We present some expressive output images from our implementation of the algorithm in GPU using RTX. We focus on visualizing the 3D flat torus.

**Flat Torus.**

Probably the most famous and easiest example of a compact 3-manifold is the *flat torus*  $\mathbb{T}^3$ . Topologically, it is obtained by gluing opposite faces of the unit cube  $[0, 1] \times [0, 1] \times [0, 1] \subset \mathbb{E}^3$ . It is easy to check that the neighborhood of each point in  $\mathbb{T}^3$  is a 3-ball of the Euclidean space. Thus  $\mathbb{T}^3$  is indeed a 3-manifold.

$\mathbb{T}^3$  is modeled by  $\mathbb{E}^3$  because it is the quotient of the Euclidean space by the group of translation spanned by  $(x, y, z) \rightarrow (x \pm 1, y, z)$ ,  $(x, y, z) \rightarrow (x, y \pm 1, z)$ , and  $(x, y, z) \rightarrow (x, y, z \pm 1)$ . Thus, the unit cube is the fundamental domain of  $\mathbb{T}^3$ .

Figure 2 presents the visualization of a scene inside the fundamental domain using the path tracing algorithm. On its left, we do not use the indirect light contribution, on the right we add the indirect light contribution considering five bounces. Note that we are using a *cornel box* with a additional window to set up our scene. The spheres inside the cube have specular, diffuse, and composed materials. The point light is above the box, close to the top window.

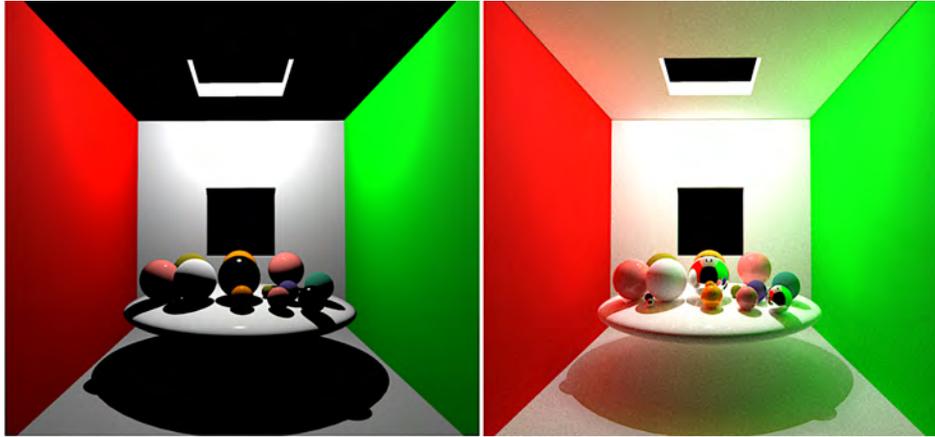


FIGURE 2. Cornell box inside the fundamental domain. Rendered considering zero/five bounces.

We now identify the boundaries of the torus fundamental domain (the cube). A ray leaving a point  $p \in \mathbb{T}^3$  in a direction  $v$  is parameterized as  $r(t) = p + t \cdot v$  in  $\mathbb{E}^3$ . For each intersection between  $r$  and a face  $F$  of the unit cube, we update  $p$  by its correspondent point  $p - n$  in the opposite face, where  $n$  is the unit vector normal to  $F$ . The ray direction  $v$  does not need to be updated. The rays in  $\mathbb{T}^3$  can return to the starting point, providing many copies of the scene. The immersive perception is  $\mathbb{E}^3$  tessellated by unit cubes: each cube contains one copy of the scene.

Figure 3 provides an immersive visualization of the 3-dimensional torus  $\mathbb{T}^3$  using the path tracer algorithm. On its top, an outside view of the embedded cornel box showing how the fundamental domain tessellates the space. On its bottom, a closer view on the specular spheres embedded in the scene to show the tessellation being reflected on their surface.

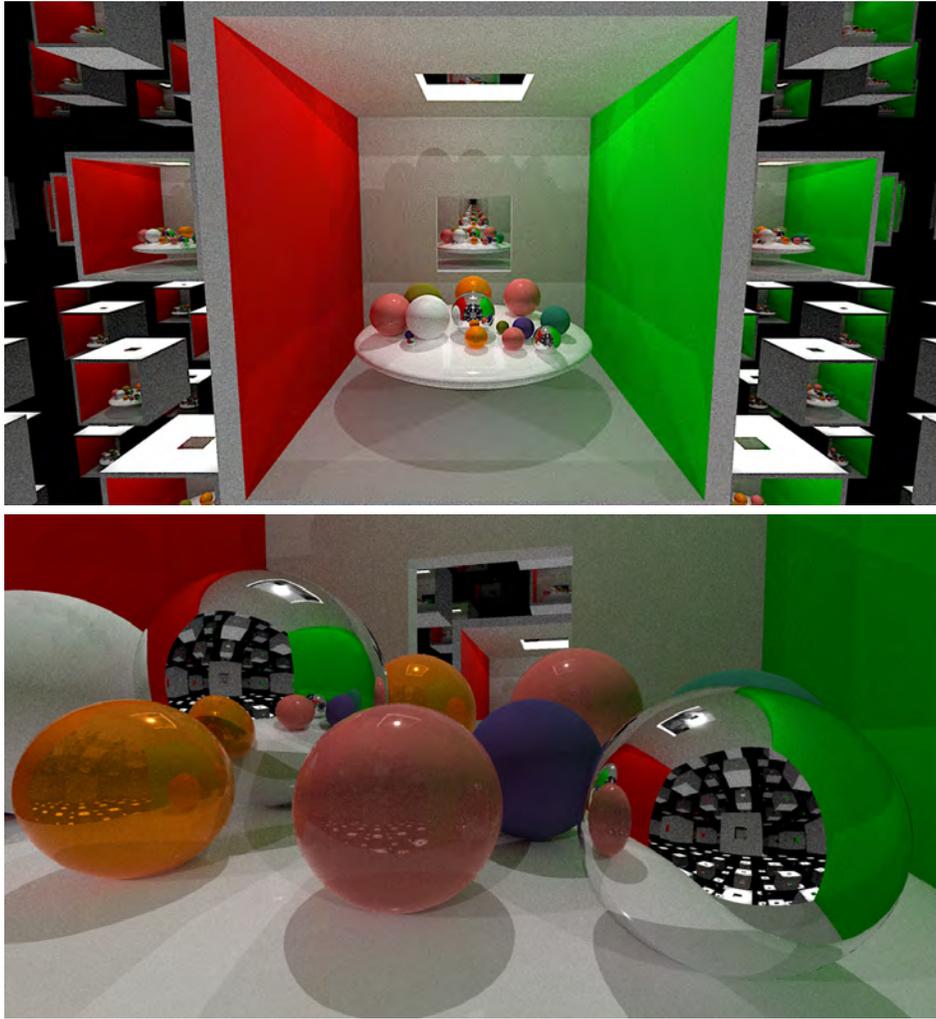


FIGURE 3. Immersive view in the 3D flat torus. The face pairing makes the rays that leave a face return from its opposite face, giving rise to many copies of the scene.

### Mirrored Dodecahedron.

For an example of a space modeled by the hyperbolic space, consider the dodecahedron embedded in  $\mathbb{H}^3$ . Let  $\Gamma$  be the group of reflections generated by the dodecahedral faces. With an appropriate scale, the dihedral angle of the dodecahedron reaches 90 degrees. The quotient  $\mathbb{H}^3/\Gamma$  is the *mirrored dodecahedral space*.  $\Gamma$  tessellates  $\mathbb{H}^3$  with dodecahedra, each edge has exactly 4 cells.

Figure 4 illustrates an immersive visualization of the mirrored dodecahedron using the reflection definition in the hyperbolic space. There are exactly three spheres, two red and a blue. We add the dodecahedron edges to highlight the hyperbolic space tessellation. The image is the inside view of the group of reflection acting on the Hyperbolic space.

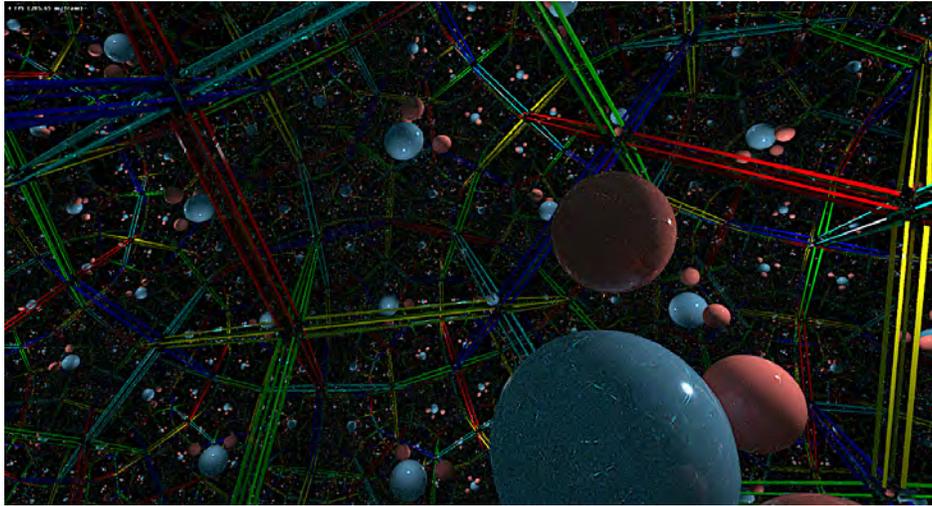


FIGURE 4. Immersive visualization of the mirrored dodecahedron. This space is obtained by considering the faces of a hyperbolic regular dodecahedron to be perfect hyperbolic mirrors.

#### REFERENCES

1. Pierre Berger, Alex Laier, and Luiz Velho, *An image-space algorithm for immersive views in 3-manifolds and orbifolds*, Visual Computer (2014).
2. James T Kajiya, *The rendering equation*, Proceedings of the 13th annual conference on Computer graphics and interactive techniques, 1986, pp. 143–150.
3. Bruno Martelli, *An introduction to geometric topology*, arXiv preprint arXiv:1610.02592 (2016).
4. Matt Pharr, Wenzel Jakob, and Greg Humphreys, *Physically based rendering: From theory to implementation*, Morgan Kaufmann, 2016.
5. W.P. Thurston, *The geometry and topology of three-manifolds*, Princeton University, 1979.
6. Luiz Velho, Tiago Novello, Vinicius Silva, and Djalma Lucio, *Visualization of non-euclidean spaces using ray tracing*, Technical Report TR-09-2019, VISGRAF Lab - IMPA, 2019.

VISGRAF LABORATORY,RIO DE JANEIRO,BRAZIL  
 E-mail address: [tiago.novello90@gmail.com](mailto:tiago.novello90@gmail.com)

VISGRAF LABORATORY,RIO DE JANEIRO,BRAZIL  
 E-mail address: [dsilva.vinicius@gmail.com](mailto:dsilva.vinicius@gmail.com)

VISGRAF LABORATORY,RIO DE JANEIRO,BRAZIL  
 E-mail address: [lvelho@impa.br](mailto:lvelho@impa.br)