# Subspace Culling for Ray–Box Intersection

ATSUSHI YOSHIMURA, Advanced Micro Devices, Inc., Japan
TAKAHIRO HARADA, Advanced Micro Devices, Inc., USA

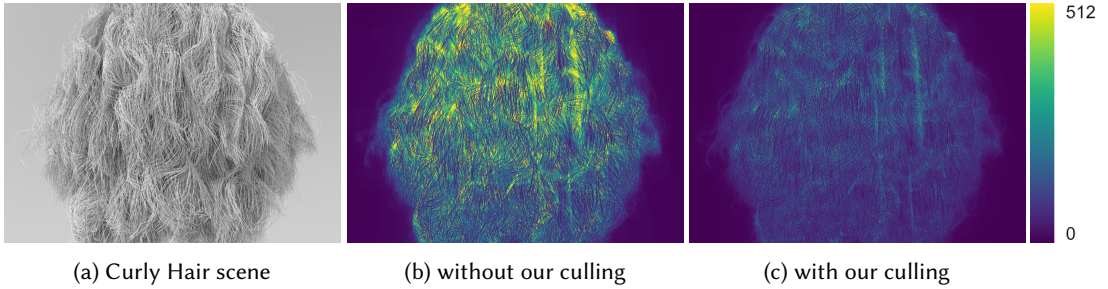| (a) Curly Hair scene | (b) without our culling | (c) with our culling |

Fig. 1. Visualization of the number of intersections on the primary rays with and without our culling, where the number of bits on the voxel data structure for each AABB is 216. The intersection between a ray and the voxels is accelerated by our look-up table-based approach. (a): The scene consists of 12.1 million triangles and almost all of the triangle is thin and tilted. (b) and (c): The intersection count of triangles and internal nodes is mapped to the color with a range of 0 - 512. Our method reduces the number of intersections by 37.5% and decreases the rendering time by 13.1% for this scene.

Ray tracing is an essential operation for realistic image synthesis. The acceleration of ray tracing has been studied for a long period of time because algorithms such as light transport simulations require a large amount of ray tracing. One of the major approaches to accelerate the intersections is to use bounding volumes for early pruning for primitives in the volume. The axis-aligned bounding box is a popular bounding volume for ray tracing because of its simplicity and efficiency. However, the conservative bounding volume may produce extra empty space in addition to its content. Especially, primitives that are thin and diagonal to the axis give false-positive hits on the box volume due to the extra space. Although more complex bounding volumes such as oriented bounding boxes may reduce more false-positive hits, they are computationally expensive. In this paper, we propose a novel culling approach to reduce false-positive hits for the bounding box by embedding a binary voxel data structure to the volume. As a ray is represented as a conservative voxel volume as well in our approach, the ray–voxel intersection is cheaply done by bitwise AND operations. Our method is applicable to hierarchical data structures such as bounding volume hierarchy (BVH). It reduces false-positive hits due to the ray–box test and reduces the number of intersections during the traversal of BVH in ray tracing. We evaluate the reduction of intersections with several scenes and show the possibility of performance improvement despite the culling overhead. We also introduce a compression approach with a lookup table for our voxel data. We show that our compressed voxel data achieves significant false-positive reductions with a small amount of memory.

CCS Concepts: • **Computing methodologies** → **Rendering**; **Ray tracing**.

Additional Key Words and Phrases: global illumination, ray tracing, voxels

Authors' addresses: Atsushi Yoshimura, Advanced Micro Devices, Inc., Japan, Atsushi.Yoshimura@amd.com; Takahiro Harada, Advanced Micro Devices, Inc., USA, Takahiro.Harada@amd.com.

## 1 INTRODUCTION

Ray tracing is a method widely used for various applications to find intersections with scene primitives. Algorithms such as path tracing rely heavily on ray tracing for photo-realistic image rendering. However, a vast number of ray queries consume significant computational resources. Several data structures and algorithms such as grid traversal [Amanatides and Woo 1987], octrees [Knoll et al. 2006], and the bounding volume hierarchy [Clark 1976] have been introduced and studied in order to accelerate ray tracing. These approaches use one or more kinds of bounding volumes to reduce the number of intersections by skipping the primitive tests that are assigned to the volume if it does not hit. The axis-aligned bounding box (AABB) is one of the most popular bounding volumes because of its simplicity and efficiency in ray traversal and data structure building. However, AABB may not be able to tightly fit the primitives such as, hair and fur, that are thin and diagonal to the axis. More spatially flexible bounding volumes, such as the oriented bounding boxes (OBB), can be used [Woop et al. 2014]. Although its culling efficiency is better, the computational cost and memory consumption are high. Instead, we propose a novel culling technique subsequent to the AABB-ray test to reduce the number of false-positive intersections produced by AABB.

Our method effectively reduces the total number of intersections by embedding a binary voxel data structure for each node in the bounding volume hierarchy (BVH), as shown in Fig. 1. A ray is represented as a binary voxel volume with the same resolution. The intersection is done by bitwise AND operations. We also introduce using a lookup table (LUT) for voxel data compression, which reduces the memory size without sacrificing a lot of culling efficiency and traversal performance.

## 2 RELATED WORK

A voxel data structure can be used for ray tracing acceleration by skipping cells that are not along the rays. Variants of a digital differential analyzer (DDA) was proposed for ray tracing [Amanatides and Woo 1987; Fujimoto et al. 1986]. Fine voxels can fit primitives tightly even if they are concave; however, the data structure requires iterative algorithms to find all intersected cells with a ray, and the number of iterations is proportional to the voxel resolutions.

Intersections between a ray and a chunk of binary voxels can be accelerated by blockwise packing [Gruen 2018]. Voxels along a ray are also packed as the same representation as the scene voxels, and the intersection is efficiently done by bitwise AND operations.

The bounding volume hierarchy [Clark 1976; Meister et al. 2021] effectively reduces the number of intersections by hierarchical culling with bounding volume such as AABB. Although AABB is the most popular choice for the bounding volume, it suffers from false positive intersections due to non-tightly fit primitives. OBB can be used for further reduction of the false positives [Woop et al. 2014]; however, the computational overhead is higher, and larger memory space is needed. Tessellating primitives into finer granularity also alleviates the false positives [Popov et al. 2009; Stich et al. 2009], but the drawback is memory consumption due to the deeper bounding volume hierarchy. Molenaar and Eisemann proposed another method using a proxy geometry represented by a hierarchical voxel data structure to optimize out-of-core rendering [Molenaar and Eisemann 2020]. Our approach is similar to theirs but our method has a smaller overhead which makes it possible to use at each step in the traversal, and it pays off for in-core rendering.
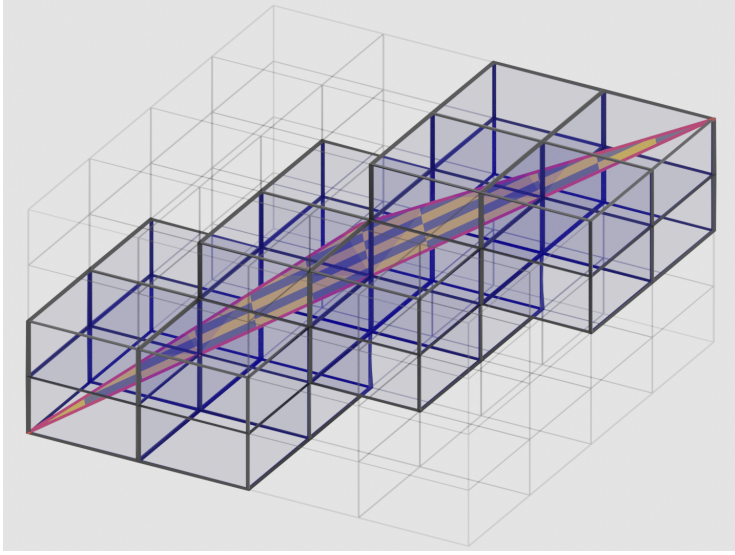
Fig. 2. A triangle bounded more tightly than AABB by $4 \times 4 \times 4$ voxels. Only 22 of 64 voxels are occupied by the triangle.

We propose a novel culling that achieves a coexistence of culling efficiency, cheap computation, and low memory cost. Our culling leverages voxels' tightness and intersection efficiency to compensate for the AABB weakness. The voxels can be compressed with LUT-based compression to minimize the memory cost.

## 3 OUR CULLING APPROACH

AABB is a simple and efficient bounding volume; however, primitives such as a long, thin, and diagonal triangle in the volume may not be tightly bounded. Our approach compensates for this loose bound due to AABB by a uniform grid and storing occupancy of a triangle as shown in Fig. 2. Since the voxel data structure can represent empty space for each cell individually, it can more tightly enclose these kinds of primitives. We build a uniform grid data structure representing spatial occupancy as binary for each AABB in a BVH, which we call an object mask. In order to determine the voxel pattern along a ray in the same grid as the object mask, AABB-ray intersections are needed, such as the DDA-based methods [Amanatides and Woo 1987; Fujimoto et al. 1986]. However, such iterative algorithms for each AABB during BVH traversal may be impractical because the grid traversal takes linear time complexity with respect to the grid resolution. Instead, we build a LUT for the overlapped voxels with a ray in the grid. Accordingly, the voxel pattern is obtained by looking up the LUT. We call the voxel pattern for a ray mask. The intersection between a ray mask and an object mask can be executed by bitwise AND operations, as shown in Fig. 3. Thus, our culling in subspace can reduce false positives due to AABB by simple arithmetic operations.

First, we describe our method with a simple case — triangles in an AABB in Sect. 3.1, then extend it to BVH, a hierarchical case, in Sect. 3.2.

### 3.1 Subspace Culling for Triangle-Ray Intersection

*3.1.1 Mask Creation.* We create an object mask and a ray mask to test the overlap. In order to build an object mask in an AABB containing a single triangle (or triangles), we conservatively voxelize
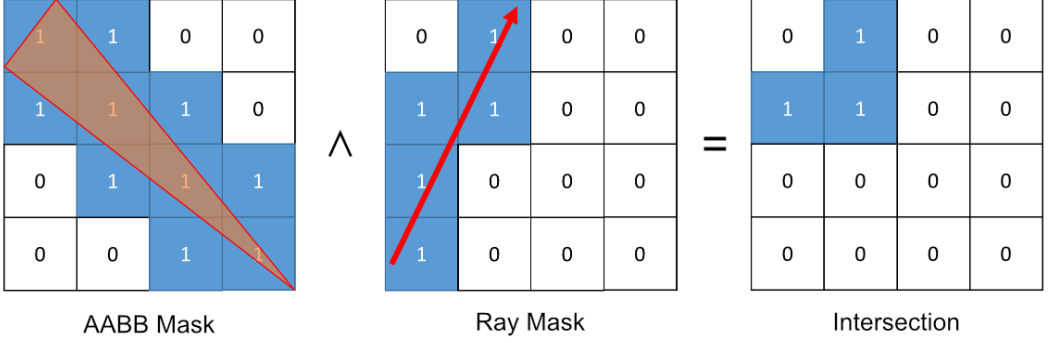
Fig. 3. An example of Ray-AABB intersection by AABB mask and ray mask by a bitwise AND operation. If all bits in the result are zero, there is no intersection.

---

**Algorithm 1:** Lookup ray mask **Inputs:** *lower, upper*: Minimum and maximum value of the AABB. $p_0$, $p_1$: The hit locations by AABB-ray intersection. *R*: The mask resolution. *rayMasks*: The LUT for a ray mask. Discretized begin and end location of the ray map to its ray mask.

---

**function** lookupRayMask( *lower, upper, $p_0$, $p_1$, R, rayMasks*)

    $extent \leftarrow upper - lower$ ;

    $beg \leftarrow \min(\lfloor \frac{R}{extent}(p_0 - lower)\rfloor, R - 1)$ ;

    $end \leftarrow \min(\lfloor \frac{R}{extent}(p_1 - lower)\rfloor, R - 1)$ ;

    **return** $rayMasks_{(beg,end)}$;

**end**

---

the triangles [Schwarz and Seidel 2010]. A chunk of binary occupancy can be densely packed as a bit sequence.

We precompute ray masks as a LUT. We use two discretized positions of AABB-ray intersections as the key of the LUT; thus, the table has six dimensions. The hit locations by AABB-ray intersection are linearly mapped to a unit cube then the indices are determined as shown in Algorithm 1 to reuse the LUT for all AABB. The two positions of AABB-ray intersections can be anywhere in the AABB in cases where the ray origin or the end of the ray segment is inside the volume, as long as the ray mask LUT includes masks corresponding to all combinations of the two positions. We use AABB-AABB Sweep intersection test [Ericson 2004] to build the LUT because the sweep represents all possible rays. Although the lookup-based ray mask is conservative, it may produce extra false-positive intersections, which we evaluate in Sect. 4.2.

*3.1.2 Culling by Ray Mask and Object Mask .* Once the ray mask LUT and object masks are created, the culling is straightforward. First, we compute the two intersection points of the ray on the AABB, which are used to fetch a ray mask from the LUT. Then, we take bitwise AND between the ray mask and the object mask. If all the bits are zero, there is no intersection between the ray and the AABB. Therefore ray-triangle intersections can be culled by the masks further after AABB culling.

## 3.2 Hierarchical Subspace Culling for BVH

Our method can be trivially applied to every node in BVH because each node can be considered a node containing all of the triangles in descendant nodes. However, calculating the occupancy with

(a) Voxel filling in the parent    (b) Two voxels are overlapped    (c) Four voxels are overlapped
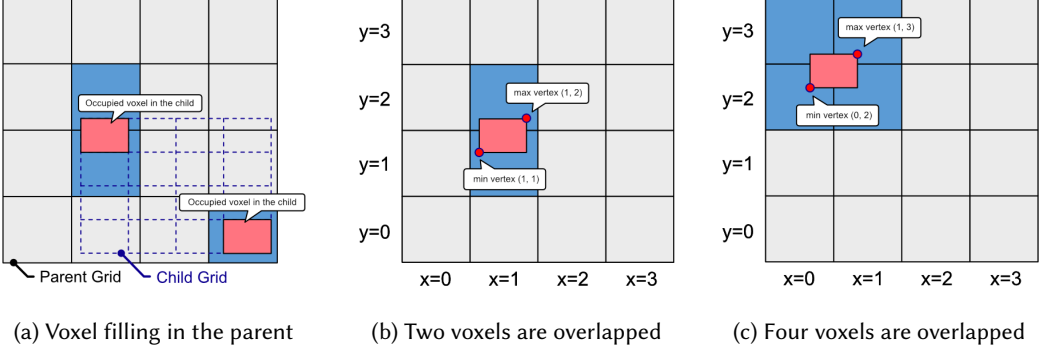
Fig. 4. The parent voxels are conservatively filled by their occupied voxels in children. Occupied voxels in the parent and occupied voxels in the child are illustrated as blue and red cells, respectively. (a): Three voxels in the parent are filled by two voxels in the child. (b) and (c): Two box-shaped filling patterns in the parent by a voxel in a child. A child voxel fills one to eight parent voxels since a child voxel size is smaller than or equal to the parent.

all primitives in a node is computationally expensive, especially for nodes in an upper level on a BVH. Thus, we propose an approximation to accelerate object mask creation dedicated to BVH. Additionally, object masks for each AABB can consume large memory. We also propose to use LUT to compress object masks in a BVH to reduce the memory overhead.

*3.2.1 Hierarchical Object Mask Creation.* As a node may contain a large number of primitives due to the hierarchical representation, voxelizing all the primitives to compute the occupancy is redundant and computationally expensive. Instead, we introduce the use of an object mask on a node as an approximation of the occupancy of its children. An occupied voxel in the child conservatively fills its parent voxels as shown Fig. 4a. Fig. 4b and Fig. 4c show two filling patterns in the parent from a child voxel. A child voxel's min and max vertices are projected to the parent grid. Their indices determine the box-shape filling pattern in the parent node. Because the pattern is finite, we precompute a six-dimensional table for the box-shaped pattern. The table maps the projected indices of the vertices to their filling pattern. Each occupied child voxel is mapped to a pattern and they are combined by bitwise OR operations to get the conservative voxels in the parent. Algorithm 2 shows an object mask creation from the approximated occupancy of its children. As the approximation may produce extra false positives, we use a parameter $L$ to control the strength of approximation, where we compute object masks from the node $L$ levels below to take tighter occupancy. As $L$ increases, the computational cost increases as we need to propagate object masks and primitives to many parents. $L = \infty$ is equivalent to the tightest object mask built from all primitives below the node while $L = 1$ limits a node to take the occupancy only from its direct children.

*3.2.2 Object Mask Compression with LUT.* The memory size of an object mask is proportional to the cube of the resolution $R$. For instance, we need to store a 64-bit value for an AABB when we use $R = 4$. In order to reduce the memory footprint of object masks, we propose the use of a LUT to compress them. Each mask data is replaced by an index of the LUT, which we call a compression LUT. The tightness of the compressed masks depends on a set of object masks we store in the compression LUT. We borrow the idea of surface area heuristic (SAH) [Goldsmith and Salmon 1987;

MacDonald and Booth 1990] to define the optimal LUT. In SAH, the cost of BVH is described as follows:

$$SAH(N) = \frac{1}{SA(N)} \left( C_T \sum_{N_i} SA(N_i) + C_I \sum_{N_l} SA(N_l) \cdot |N_l| \right), \tag{1}$$

where $N$ is the root node, $SA(N)$ is the surface area of the bounding box on the node, $N_i$ and $N_l$ are inner nodes and leaf nodes, $C_T$ and $C_I$ are intersection cost of an inner node and a primitive, $|N_l|$ is the number of primitives in the leaf.

We assume the number of occupancy in the object mask is proportional to the probability of the intersecting ray. Then the SAH can be extended to take our culling into account as follows:

$$SAH_{masked}(N) = \frac{1}{SA(N)} \left( C_T \sum_{N_i} SA(N_i) \frac{O_m(N_i)}{R^3} + C_I \sum_{N_l} SA(N_l) \frac{O_m(N_l)}{R^3} \cdot |N_l| \right), \tag{2}$$

where $O_m(x)$ is the number of occlusions in the object mask in a node $x$. As a compression LUT has a limited number of elements, the LUT that minimizes Eq. 2 is the optimal table.

However, it is prohibitively expensive to find the optimal LUT because of the vast search space. Thus, we first choose the object masks from the original masks for a compression LUT. We use $SA(N)\frac{O_m(N)}{R^3}$ as the importance of its object mask for creating a sub-optimal LUT. We take the top important object masks as a compression LUT. Second, we find the conservative and the tightest object mask from the LUT for each mask that is not found in the LUT.

Note that the intersection of the LUT for ray mask and compression LUT can be obtained by a precomputed two-dimensional bit table. The table key consists of an index of the ray mask and an index of the compression LUT. The table values are the intersection result represented as a bit. Accordingly, fetching a ray mask and an object mask from each LUT and bitwise AND operations can be replaced by fetching an intersection result as a bit from the precomputed bit table. This replacement can improve the performance of the intersection if the bit fetching is cheaper despite memory access to random locations in the large bit table.

*3.2.3 Search Object Mask in the LUT .* Some object masks may not be in the compression LUT when the number of LUT elements is less than object masks in a BVH. We need to select an alternative mask from compression LUT for such missing masks. The alternative mask needs to be conservative; otherwise, it causes invalid culling. The alternative mask is preferred to be a tighter one in the conservative masks to make Eq. 2 smaller. Accordingly, we extract masks that have all set bits of the missing object mask out of compression LUT and choose the mask with the smallest number of set bits. The brute-force algorithm can do this; however, we propose a search algorithm with which we can search for the best mask in table look-ups and a few arithmetic operations. The algorithm consists of conservative mask extraction and searching for the tightest one. The former is done using precomputed tables, and the latter is done by presorting the compression LUT by their set bit count.

Since the set bits in the object mask indicate the requirements, we precompute the list of masks that satisfy the requirements in the compression LUT against each bit pattern. The list of the masks is stored as a bit sequence where each bit represents whether the mask corresponding to its bit location satisfies the requirements. However, the bit pattern is $2^{R^3}$, which is impractically large to compute. In order to reduce the size of the precomputed table, we separate the bit pattern into smaller chunks, where each chunk has $b$ bits, and each precomputed table has $2^b$ elements. We call the table requirement LUT. Instead of building a single massive table, we create $\lceil \frac{R^3}{b} \rceil$ requirement LUTs built from compression LUT as we prepare a table for all bit chunks. A list of masks out of

---

**Algorithm 2:** An object mask creation from the approximated occupancy. The indices of the parent voxel grid are calculated individually per axis to reuse the result in the loop for each voxel of the child mask. **Inputs and Symbols:** *mask*: AABB mask to update its occupancy. *pLower, pUpper*: Minimum and maximum value of the parent AABB. *childMask*: Approximated occupancy of the child. *cLower, cUpper*: Minimum and maximum value of the child AABB. *R*: Mask resolution. *fillingPatternTable: A table maps minimum and maximum voxel indices in a parent grid to a box-shaped voxel pattern. ⊕: A bitwise OR.*

---

**function** fillByApproximatedOccupancy( *mask, pLower, pUpper, childMask, cLower, cUpper, R, fillingPatternTable*)

> $cExtent \leftarrow cUpper - cLower$ ;
> $pExtent \leftarrow pUpper - pLower$ ;
> **for** $i \leftarrow 0$ *to R* **do**
>> $border \leftarrow cLower + \frac{i}{R}cExtent$ ;
>> $indexParent \leftarrow \lfloor R\frac{border-pLower}{pExtent} \rfloor$ ;
>> $xs_i, ys_i, zs_i \leftarrow \min(indexParent, R-1)$;
>
> **end**
>
> // A Loop for each voxel of the child mask.
> **for** $x \leftarrow 0$ *to R - 1* **do**
>> **for** $y \leftarrow 0$ *to R - 1* **do**
>>> **for** $z \leftarrow 0$ *to R - 1* **do**
>>>> **if** $childMask_{xyz} = 0$ **then**
>>>>> **continue**
>>>>
>>>> **end**
>>>> $index_{min} \leftarrow xs_x, ys_y, zs_z$;
>>>> $index_{max} \leftarrow xs_{x+1}, ys_{y+1}, zs_{z+1}$;
>>>> $mask \leftarrow mask \oplus fillingPatternTable_{(index_{min},index_{max})}$ ;
>>>
>>> **end**
>>
>> **end**
>
> **end**

**end**

---

the requirement LUTs can be combined into a bit sequence by bitwise AND operations, where the combined bit sequence indicates the list of masks that satisfy all of the bit requirements of an object mask as shown in Fig. 5. The object mask can be conservatively replaced by any mask in the list.

Once the conservative mask list is obtained, the mask with the lowest set bit count in the list is the tightest. Instead of the linear search, we use the least-significant bit (LSB) location to find the tightest mask by sorting the masks in the compression LUT order by their number of set bits before the requirement LUTs are built. Algorithm 3 shows our mask search algorithm.

*3.2.4 BVH Node Structure.* Fig. 6 shows memory layout visualization of 4-wide BVH nodes with our object masks as examples, where $R = 4$, the number of the compression LUT element is 256, Fig. 6a is without the compression, and Fig. 6b is with the compression. Since the number of bits in the object mask is $R^3$, each mask in the uncompressed node is 8 bytes. The index of a mask in the compression LUT can be stored in 1 byte, which is $\frac{1}{8}$ of the uncompressed mask. The size of the
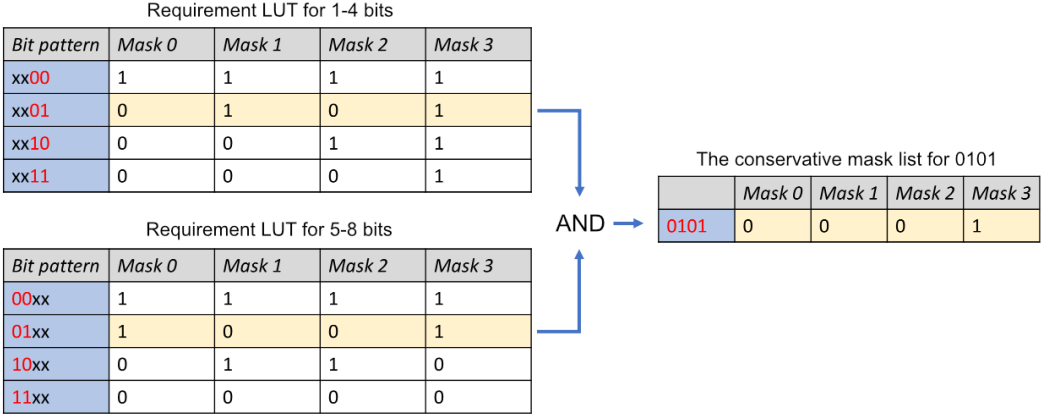
**Requirement LUT for 1-4 bits**

| Bit pattern | Mask 0 | Mask 1 | Mask 2 | Mask 3 |
|---|---|---|---|---|
| xx00 | 1 | 1 | 1 | 1 |
| xx01 | 0 | 1 | 0 | 1 |
| xx10 | 0 | 0 | 1 | 1 |
| xx11 | 0 | 0 | 0 | 1 |

**The conservative mask list for 0101**

| | Mask 0 | Mask 1 | Mask 2 | Mask 3 |
|---|---|---|---|---|
| 0101 | 0 | 0 | 0 | 1 |

AND →

**Requirement LUT for 5-8 bits**

| Bit pattern | Mask 0 | Mask 1 | Mask 2 | Mask 3 |
|---|---|---|---|---|
| 00xx | 1 | 1 | 1 | 1 |
| 01xx | 1 | 0 | 0 | 1 |
| 10xx | 0 | 1 | 1 | 0 |
| 11xx | 0 | 0 | 0 | 0 |

Fig. 5. An example of requirement LUTs with $b = 2$. A bit pattern of an object mask is separated into $b$ bits elements, and a mask list that corresponds to its bit pattern from each requirement LUT is combined by bitwise AND operations to get the conservative mask list for the object mask to search.

---

**Algorithm 3:** Search the optimal object mask in compression LUT. **Inputs, Functions, and Notations:** *mask*: An object mask to search for. $R$: The mask resolution. $b$: The number of bits in a chunk that corresponds to the requirement table. *requirementTables*: The precomputed requirement tables. $LSB(x)$: Calculate least significant bit of x. $\wedge$: A bitwise AND.

---

**function** indexOfOptimalMask( *mask, R, b, requirementTables*)
    $nBatch \leftarrow \lceil \frac{R^3}{b} \rceil$ ;
    **for** $i \leftarrow 0$ *to* $nBatch - 1$ **do**
        $bits \leftarrow$ read $b$ bits from *mask* at $(i \cdot b)$-th bit ;
        $requirement_i \leftarrow requirementTables_{(i,bits)}$ ;
    **end**
    $conservativeList \leftarrow requirement_0$ ;
    **for** $i \leftarrow 1$ *to* $nBatch - 1$ **do**
        $conservativeList \leftarrow conservativeList \wedge requirement_i$;
    **end**
    **return** LSB ( $conservativeList$ );
**end**

---

uncompressed node is 144 bytes compared to 112 bytes without masks. The size of a node with compression is 116 bytes.

*3.2.5 Traversal Algorithm.* Our subspace culling algorithm can be added to an existing regular BVH traversal algorithm. A pseudocode is shown in Algorithm 4 where the addition of the logic for the proposed method is highlighted in red. Our method reuses the results of the AABB-ray intersection.
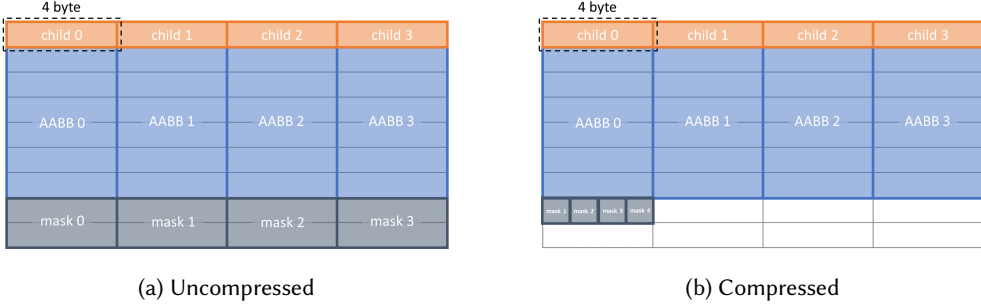
(a) Uncompressed          (b) Compressed

Fig. 6. Visualization of memory layout for 4-wide BVH nodes with our object masks, where $R = 4$. Each rectangle cell indicates 4 bytes, and each square cell indicates 1 byte. (a): A BVH node without compression. (b): A BVH node with our compression by the LUT.

---

**Algorithm 4:** BVH traversal with our culling. The red highlighted part is our culling embedded in a traditional BVH traversal. **Inputs, Functions, and Notations:** *root*: A root node of the BVH. *ray*: A ray to find intersections to the BVH. *R*: The mask resolution. *rayMasks*: The LUT for a ray mask. Discretized begin and end location of the ray map to its ray mask. *lookupRayMask*: LUT-based ray mask. Creation of the LUT is described in Algorithm 1

---

**function** traverse( *root, ray, R, rayMasks*)
    push (*root*);
    **while** *True*
        *node* ← pop () ;
        **if** *node is empty* **then break** ;
        **if** *node is a leaf* **then**
            Find an intersection *ray* and triangles in *node* ;
            **continue**;
        **end**
        *hits* ← find intersections AABBs in *node* and *ray* ;
        **foreach** $hit_i \in hits$ **do**
            *AABB, objectMask, ps* ← get an AABB, an object mask, and intersections of $hit_i$;
            *rayMask* ← lookupRayMask ( $AABB_{lower}, AABB_{upper}, ps_0, ps_1, R, rayMasks$ );
            **if** *objectMask* ∧ *rayMask is not zero* **then**
                *child* ← get an child node that corresponds $hit_i$;
                push (*child*);
            **end**
        **end**
    **end**
**end**

---

## 4 EVALUATION

We evaluate the culling efficiency in path tracing with several scenes shown in Fig. 7. All the measurements are done with $960 \times 540$ resolution and 16 samples per pixel. We count the number

(a) Bedroom (462.8 K tris)

(b) San Miguel (9.9 M tris)

(c) Ninja (1.3 M tris)

(d) Bistro (2.8 M tris)

(e) Classroom (606.1 K tris)

(f) Hairball (2.8 M tris)

(g) Curly Hair (12.1 M tris)

(h) Straight Hair (7.3 M tris)

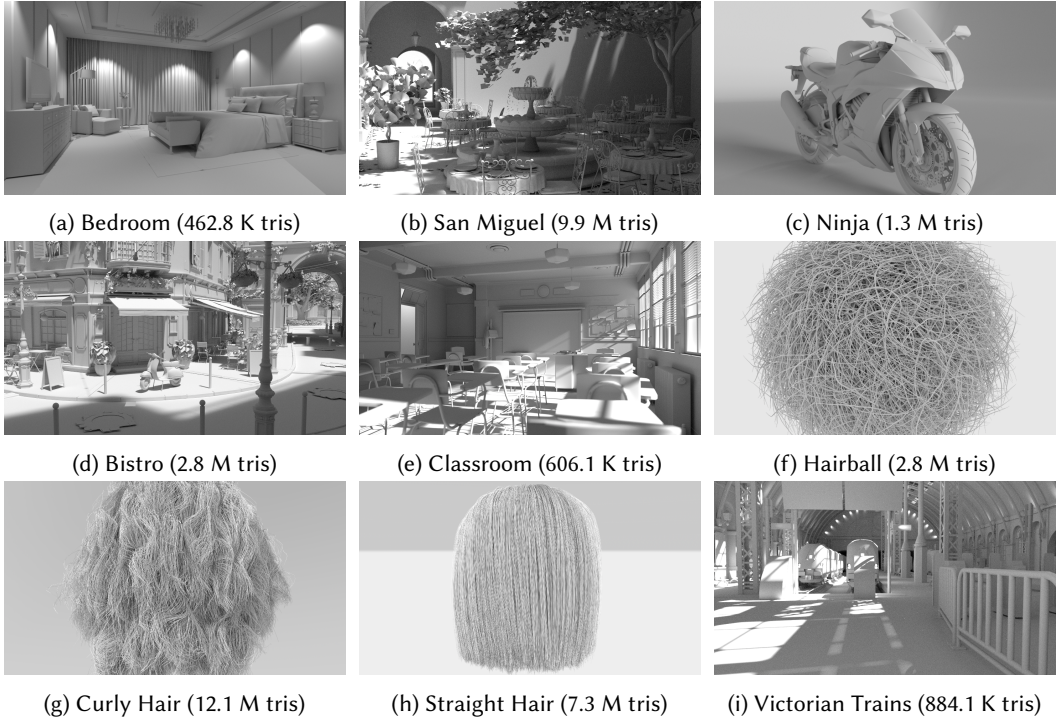(i) Victorian Trains (884.1 K tris)

Fig. 7. Test scenes used to evaluate our subspace culling.

of intersections of AABB-ray and triangle-ray to evaluate the culling efficiency. As we use 4-wide BVH for all measurements in this paper, we treat four AABB-ray intersections and a triangle-ray intersection as a measurement unit. The BVH for each scene is built by binned SAH BVH construction algorithm [Meister et al. 2021]. Note that our method does not assume any branching factor or BVH construction algorithm; therefore, it applies to a BVH with any of them. The maximum culling efficiency with our method is achieved by ray masks created with iterative grid traversal and object masks built without the approximation where an object mask is calculated by all of the descendant primitives in the BVH. First, we show the maximum culling efficiency. Next, we evaluate the increase of intersections and their rendering performance with the lookup table optimization of the ray mask. The trade-off between the culling efficiency and the performance of object mask building with approximated object masks is also shown. Finally, we show the culling efficiency of the compressed masks and their object mask-finding performance. All measurements were done with an AMD Ryzen™9 5950X CPU.

## 4.1 Ideal Culling

We use grid traversal proposed by Amanatides et al. for the maximum culling efficiency [Amanatides and Woo 1987]. "Ideal Culling" column in Table 1 shows the ratio of intersections compared to without our culling. The ideal intersection ratio is 75.4% on average, 56.5% on the best with $R = 4$, and 65.3% on average, 42.5% on the best with $R = 6$. Significant intersection reductions are observed in the "Curly Hair" and "Straight Hair" scene compared to the others. It emphasizes our method is suitable for thin and tilted geometry. Despite the reduction of intersections, it is difficult to pay off the culling overhead due to the ray mask creation by grid traversal as shown in Table 2.

| Ratio of intersections with resolution $R = 4$ | | | | | |
|---|---|---|---|---|---|
| Scene | No Culling | Ideal Culling | Ray mask LUT $R_{ray} = R$ | Ray mask LUT $R_{ray} = 2R$ | Compressed with LUT |
| Bedroom | 100% | 89.3% | 90.0% | 89.8% | 90.7% |
| San Miguel | 100% | 79.3% | 80.5% | 80.2% | 83.2% |
| Ninja | 100% | 84.2% | 85.5% | 85.2% | 90.0% |
| Bistro | 100% | 72.3% | 73.8% | 73.4% | 78.2% |
| Classroom | 100% | 87.4% | 88.1% | 87.9% | 88.3% |
| Hairball | 100% | 67.0% | 71.5% | 70.6% | 73.2% |
| Curly Hair | 100% | 57.5% | 62.5% | 61.5% | 66.3% |
| Straight Hair | 100% | **56.5%** | **61.4%** | **60.3%** | **62.1%** |
| VictorianTrains | 100% | 85.1% | 86.2% | 85.9% | 86.9% |
| *Average* | 100% | 75.4% | 77.7% | 77.2% | 79.9% |
| Ratio of intersections with resolution $R = 6$ | | | | | |
| Bedroom | 100% | 83.8% | 84.9% | 84.5% | 86.9% |
| San Miguel | 100% | 67.9% | 70.6% | 69.7% | 76.5% |
| Ninja | 100% | 76.5% | 78.8% | 78.0% | 87.7% |
| Bistro | 100% | 59.1% | 61.8% | 61.0% | 71.5% |
| Classroom | 100% | 83.5% | 83.3% | 83.7% | 85.1% |
| Hairball | 100% | 54.4% | 61.1% | 58.8% | 74.3% |
| Curly Hair | 100% | 43.3% | 50.6% | 48.1% | 60.4% |
| Straight Hair | 100% | **42.5%** | **49.3%** | **47.0%** | **53.6%** |
| VictorianTrains | 100% | 77.0% | 78.7% | 78.1% | 81.1% |
| *Average* | 100% | 65.3% | 68.8% | 67.6% | 75.2% |

Table 1. Ratio of AABB-Ray and AABB-Triangle intersections with and without our culling. Ideal culling case, ray mask with LUT approximation, and compressed cases are shown. The underlined numbers are the best ratio in the scenes.

## 4.2 Ray Mask Creation with LUT

Although the LUT for ray masks can be used to avoid the expensive grid traversal, extra false positives are produced due to the use of the LUT. We evaluate the increase of intersections with ray mask LUT and the rendering performance. The LUT resolution–$R_{ray}$ can be independent of mask resolution–$R$. However, a LUT cell can overlap with more than one cell in an AABB mask if $R_{ray}$ is not multiple of $R$ or $R_{ray}$ is less than $R$. These misaligned grids reduce the method effectiveness because it produces false positives which could be culled. Thus, we use $R_{ray} = R$ and $R_{ray} = 2R$ for the evaluation. The creation of ray mask LUT with $R = 4$ and $R = 6$ take 1.1 milliseconds and 29.6 milliseconds, respectively, if $R_{ray} = R$, while they take 54.2 milliseconds and 1791.7 milliseconds if $R_{ray} = 2R$. The two intersection ratios and performance measurements are also shown in Table 1 and Table 2, respectively. Although the intersection ratio of ray mask LUT $R_{ray} = R, 2R$ is slightly worse than the ideal case, all scenes get faster than the ideal culling. Especially, over 10% render time reductions from the baseline are observed in "Curly Hair" and "Straight Hair" scenes with $R_{ray} = R$. $R_{ray} = 2R$ case is slower than $R_{ray} = R$ despite of the better tightness. This can be explained latency penalty of the memory access to random locations in the larger size of the LUT.

## 4.3 Approximated Object Mask

We evaluate the approximation of the object mask construction in a BVH with different parameters where $L = 1, 2, 3, 4, 5$. As the tightness and its computational cost depend on the parameter $L$, we measure the intersection and the creation time of the mask. Fig. 8 shows the trade-off between

| Relative rendering time with resolution $R = 4$ | | | | |
|---|---|---|---|---|
| Scene | No Culling | Ideal Culling | Ray mask LUT $R_{ray} = R$ | Ray mask LUT $R_{ray} = 2R$ |
| Bedroom | 100% ( 73.89 seconds ) | 117.2% | 103.6% | 103.0% |
| San Miguel | 100% ( 180.03 seconds ) | 138.0% | 97.6% | 99.3% |
| Ninja | 100% ( 55.19 seconds ) | __113.7__% | 101.0% | 101.2% |
| Bistro | 100% ( 97.01 seconds ) | 124.8% | 94.9% | 97.1% |
| Classroom | 100% ( 123.31 seconds ) | 120.2% | 106.0% | 102.9% |
| Hairball | 100% ( 54.11 seconds ) | 132.1% | 97.9% | 98.0% |
| Curly Hair | 100% ( 82.74 seconds ) | 126.7% | __86.9__% | __89.2__% |
| Straight Hair | 100% ( 54.34 seconds ) | 120.0% | 88.0% | 91.1% |
| VictorianTrains | 100% ( 129.19 seconds ) | 124.4% | 104.3% | 107.2% |
| *Average* | 100% ( 94.42 seconds ) | 124.1% | 97.8% | 98.8% |
| Relative rendering time with resolution $R = 6$ | | | | |
| Bedroom | 100% ( 73.89 seconds ) | 119.7% | 106.5% | 107.9% |
| San Miguel | 100% ( 180.03 seconds ) | 140.0% | 98.2% | 108.4% |
| Ninja | 100% ( 55.19 seconds ) | 118.1% | 104.5% | 110.3% |
| Bistro | 100% ( 97.01 seconds ) | 127.1% | 97.3% | 105.7% |
| Classroom | 100% ( 123.31 seconds ) | 124.6% | 107.1% | 110.8% |
| Hairball | 100% ( 54.11 seconds ) | 132.5% | 102.0% | 110.9% |
| Curly Hair | 100% ( 82.74 seconds ) | 120.2% | 87.1% | 95.7% |
| Straight Hair | 100% ( 54.34 seconds ) | __114.6__% | __85.6__% | __90.5__% |
| VictorianTrains | 100% ( 129.19 seconds ) | 128.4% | 110.0% | 115.8% |
| *Average* | 100% ( 94.42 seconds ) | 125.0% | 99.8% | 106.2% |

Table 2. Rendering time comparison for ray mask LUT. The ratios are relative rendering times with respect to without culling. The underlined numbers are the best performance in the scenes.

the intersection ratio and the mask creation performance. However, proper $L$ depends on the performance gain from the reduction and the amount of ray-tracing cost. As the reduction of intersection starts saturated around $L = 3$ while the mask creation time increase near to linear with respect to $L$ on almost all scenes, $L = 3$ could be used as a default parameter before fine-tuning.

### 4.4 Compressed Object Mask

As the bit pattern of the object mask can be alternated by another object mask as long as it is conservative, the number of the LUT elements can be arbitrary. We evaluate the compression with 256 for the LUT size for simplicity of the compression and decompression. The compression rate is fixed, which is 1:8 for $R = 4$ and 1:27 for $R = 6$. Table 1 shows the intersection ratio as "Compressed with LUT" column. Despite the compressed size being the same, $R = 6$ case produces better culling ratios with compression. 24.8% on average, 46.4% on maximum was reduced by just 8bits mask data per AABB in $R = 6$ case.

We also evaluate the performance of the object mask finding. We use $b = 8$ for mask search for efficient bit extraction from the target object mask and keeping the table size small. Table 3 shows a performance comparison between the brute-force search and our search. The latter only takes 9.1%−11.1% with $R = 4$, 10.7%−17.2% with $R = 6$, compared with the former.

## 5 CONCLUSIONS AND FUTURE WORK

We proposed a novel culling technique using ray and object masks that can be embedded in a BVH node. Our ray LUT-based approach reduces the number of intersections by 38.6% with $R = 4$ and

(a) Ratio of intersections with $R = 4$



(b) Ratio of intersections with $R = 6$



(c) Object mask creation time with $R = 4$

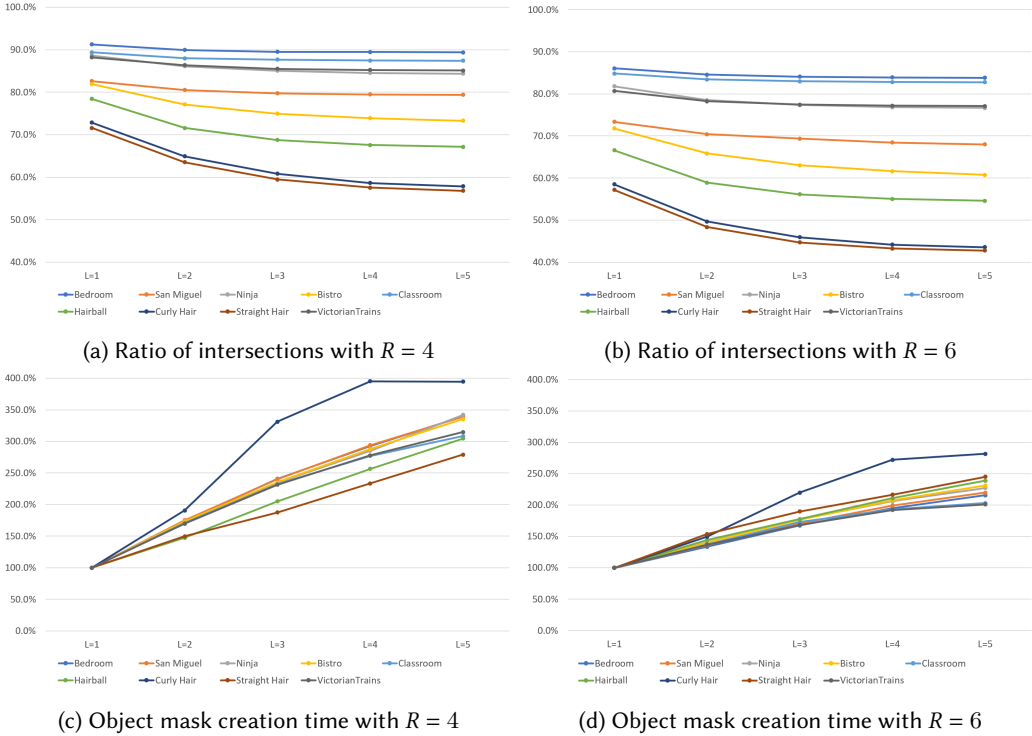

(d) Object mask creation time with $R = 6$

Fig. 8. (a) and (b): Ratio of intersections with object mask approximation with parameter $L = 1, 2, 3, 4, 5$ and the resolution $R = 4, 6$. (c) and (d): Object mask relative creation time with the approximation with parameter $L = 1, 2, 3, 4, 5$ and the resolution $R = 4, 6$. The baseline of the ratio is $L = 1$,

50.7% with $R = 6$ as the maximums in our experiment. Additionally, the traversal performance with the ray LUT approach shows performance improvement in certain scenes where AABB causes too many false positive intersections due to the loose fit. Approximated object mask provides a trade-off for object mask construction cost and culling efficiency. We also proposed compression of object masks using LUT without sacrificing much culling efficiency.

Despite our algorithm's simplicity and the intersection reduction, it is difficult to achieve a better performance in all scenes, as we showed in Sec. 4, even though our computational overhead is small. A dedicated hardware implementation may change the algorithm trade-off between computational cost and culling efficiency. Optimal LUT creation for object mask compression in a practical time has yet to be established. Fast and optimal compression LUT generation is important for the coexistence of culling efficiency, cheap computation, and low memory cost.

As each object mask in a BVH is affected by all primitives in their children's primitives, animated primitives enforce updating parent object masks. Efficient update approaches in dynamic scenes are left for future work. Another future work is an efficient object mask creation on the GPU and performance evaluation of the proposed method on the GPU.

# 6 ACKNOWLEDGMENT

| Mask finding time $R = 4$ | | | |
|---|---|---|---|
| Scene | Naive (milliseconds) | Our (milliseconds) | Our (Relative) |
| Bedroom | 216.6 | 20.3 | 9.4% |
| San Miguel | 4923.2 | 446.6 | 9.1% |
| Ninja | 569.6 | 63.1 | 11.1% |
| Bistro | 1295.2 | 128.5 | 9.9% |
| Classroom | 279.7 | 26.4 | 9.4% |
| Hairball | 1396.5 | 132.2 | 9.5% |
| Curly Hair | 5325.4 | 543.0 | 10.2% |
| Straight Hair | 3132.5 | 326.7 | 10.4% |
| VictorianTrains | 409.8 | 37.4 | 9.1% |
| Mask finding time $R = 6$ | | | |
| Bedroom | 371.2 | 46.8 | 12.6% |
| San Miguel | 9181.5 | 983.7 | 10.7% |
| Ninja | 1239.8 | 139.1 | 11.2% |
| Bistro | 2646.1 | 305.9 | 11.6% |
| Classroom | 515.4 | 60.9 | 11.8% |
| Hairball | 2039.0 | 306.8 | 15.0% |
| Curly Hair | 8768.3 | 1251.3 | 14.3% |
| Straight Hair | 4522.9 | 777.2 | 17.2% |
| VictorianTrains | 674.5 | 95.0 | 14.1% |

Table 3. Time of searching the conservative and tightest object mask in a table for compression.

## REFERENCES

John Amanatides and Andrew Woo. 1987. A Fast Voxel Traversal Algorithm for Ray Tracing. In *EG 1987-Technical Papers*. Eurographics Association. https://doi.org/10.2312/egtp.19871000

James H. Clark. 1976. Hierarchical Geometric Models for Visible-Surface Algorithms. 10, 2 (jul 1976), 267. https://doi.org/10.1145/965143.563323

Christer Ericson. 2004. *Real-Time Collision Detection*. CRC Press, Inc., USA.

Akira Fujimoto, Takayuki Tanaka, and Kansei Iwata. 1986. ARTS: Accelerated Ray-Tracing System. *IEEE Computer Graphics and Applications* 6, 4 (1986), 16–26. https://doi.org/10.1109/MCG.1986.276715

Jeffrey Goldsmith and John Salmon. 1987. Automatic Creation of Object Hierarchies for Ray Tracing. *IEEE Computer Graphics and Applications* 7, 5 (1987), 14–20. https://doi.org/10.1109/MCG.1987.276983

Holger Gruen. 2018. Block-Wise Linear Binary Grids for Fast Ray-Casting Operations. In *GPU Pro 360 Guide to Rendering* (1st ed.). A. K. Peters, Ltd., USA, Chapter 16.

Aaron Knoll, Ingo Wald, Steven Parker, and Charles Hansen. 2006. Interactive Isosurface Ray Tracing of Large Octree Volumes. In *2006 IEEE Symposium on Interactive Ray Tracing*. 115–124. https://doi.org/10.1109/RT.2006.280222

David J. MacDonald and Kellogg S. Booth. 1990. Heuristics for Ray Tracing Using Space Subdivision. *Vis. Comput.* 6, 3 (may 1990), 153–166. https://doi.org/10.1007/BF01911006

Daniel Meister, Shinji Ogaki, Carsten Benthin, Michael J. Doyle, Michael Guthe, and Jiří Bittner. 2021. A Survey on Bounding Volume Hierarchies for Ray Tracing. *Computer Graphics Forum* 40, 2 (2021), 683–712. https://doi.org/10.1111/cgf.142662 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.142662

Mathijs Molenaar and Elmar Eisemann. 2020. Conservative Ray Batching using Geometry Proxies. In *Eurographics 2020 - Short Papers*, Alexander Wilkie and Francesco Banterle (Eds.). Eurographics, The Eurographics Association. https://doi.org/10.2312/egs.20201006 DOI: 10.2312/egs.20201006.

Stefan Popov, Iliyan Georgiev, Rossen Dimov, and Philipp Slusallek. 2009. Object partitioning considered harmful: space subdivision for BVHs. In *HPG '09: Proceedings of the 1st ACM conference on High Performance Graphics* (New Orleans,

Louisiana). ACM, New York, NY, USA, 15–22. https://doi.org/10.1145/1572769.1572772

Michael Schwarz and Hans-Peter Seidel. 2010. Fast Parallel Surface and Solid Voxelization on GPUs. 29, 6, Article 179 (dec 2010), 10 pages. https://doi.org/10.1145/1882261.1866201

Martin Stich, Heiko Friedrich, and Andreas Dietrich. 2009. Spatial Splits in Bounding Volume Hierarchies. In *Proceedings of the Conference on High Performance Graphics 2009* (New Orleans, Louisiana) *(HPG '09)*. Association for Computing Machinery, New York, NY, USA, 7––13. https://doi.org/10.1145/1572769.1572771

Sven Woop, Carsten Benthin, Ingo Wald, Gregory S. Johnson, and Eric Tabellion. 2014. Exploiting Local Orientation Similarity for Efficient Ray Traversal of Hair and Fur. In *Eurographics/ ACM SIGGRAPH Symposium on High Performance Graphics*, Ingo Wald and Jonathan Ragan-Kelley (Eds.). The Eurographics Association. https://doi.org/10.2312/hpg.20141092