

DOBB-BVH: Efficient Ray Traversal by Transforming Wide BVHs into Oriented Bounding Box Trees using Discrete Rotations

Michael A. Kern¹  Alain Galvan¹  David Oldcorn¹ Daniel Skinner¹  Rohan Mehalwal¹  Leo Reyes Lozano¹
Matthäus G. Chajdas² 

¹Advanced Micro Devices, Inc.

²Intel Corporation




Figure 1: Performance comparison and visualization of the lowest interior hierarchy nodes and their child bounding volumes in our DOBB-BVH, shown alongside its corresponding AABB-based BVH and the original triangle mesh.

Abstract

Oriented bounding box (OBB) bounding volume hierarchies offer a more precise fit than axis-aligned bounding box hierarchies in scenarios with thin elongated and arbitrarily rotated geometry, enhancing intersection test performance in ray tracing. However, determining optimally oriented bounding boxes can be computationally expensive and have high memory requirements. Recent research has shown that pre-built hierarchies can be efficiently converted to OBB hierarchies on the GPU in a bottom-up pass, yielding significant ray tracing traversal improvements. In this paper, we introduce a novel OBB construction technique where all internal node children share a consistent OBB transform, chosen from a fixed set of discrete quantized rotations. This allows for efficient encoding and reduces the computational complexity of OBB transformations. We further extend our approach to hierarchies with multiple children per node by leveraging Discrete Orientation Polytopes (k-DOPs), demonstrating improvements in traversal performance while limiting the build time impact for real-time applications. Our method is applied as a post-processing step, integrating seamlessly into existing hierarchy construction pipelines. Despite a 12.6% increase in build time, our experimental results demonstrate an average improvement of 18.5% in primary, 32.4% in secondary rays, and maximum gain of 65% in ray intersection performance, highlighting its potential for advancing real-time applications.

CCS Concepts

• **Computing methodologies** → Ray Tracing; Massively parallel algorithms;

1. Introduction

In the computer graphics and gaming industry, recent advances in GPU hardware have enabled developers to synthesize photorealistic images and simulate light transport in real time. However, achieving this requires highly efficient intersection tests and occlusion queries. In ray tracing, bounding volume hierarchies (BVHs) with axis-aligned bounding boxes (AABBs) remain the predominant acceleration structure [MOB*21] to efficiently determine the

closest or arbitrary intersections of light rays with scene geometry. Although ray-AABB intersection tests are optimized in modern hardware [AMD25, NVI25, INT24], they often fail to tightly fit certain geometry, particularly thin elongated geometries not aligned with the three principal axes, such as hair strands or foliage. This inefficiency reduces the effectiveness of box culling and traversal performance. To address this issue, prior research in collision detection and ray tracing has explored alternative bounding volumes that

provide tighter fits and improved hit detection performance. These include spheres, oriented bounding boxes (OBBs), capsules, and discrete oriented polytopes (k -DOPs) [SE02, Eri04]. For elongated, arbitrarily rotated geometry, spheres tend to overestimate bounding volumes, while k -DOPs introduce higher computational complexity and memory requirements. In contrast, orientation-aware bounding volumes, such as OBBs, offer a compelling alternative to AABBs, as they can closely fit arbitrarily oriented geometry with only a modest increase in computational cost and intersection complexity. Previous studies have shown that for structures like hair strands, cables, or branches following a dominant axis in 3D space, OBBs improve hit inference performance [WBW*14] in collision detection [RW80] or view frustum culling [SGM96].

Nevertheless, the adoption of OBBs in real-time ray tracing applications has been hindered by three major challenges: (a) fast and robust OBB construction, (b) low memory representation of OBBs within hierarchical nodes, and (c) efficient ray-OBB intersection and traversal. Recent research has made significant progress in fast and robust OBB construction [VEPG23, KB24], demonstrating that OBB bounding volume hierarchies (BVH) can be generated in milliseconds by converting AABB or k -DOP hierarchies to OBBs as a post-processing step. However, existing methods and studies rely on binary BVHs, do not impose constraints on BVH memory usage, and require non-SIMD-friendly ray-OBB intersection tests for each child node in the hierarchy. In this paper, we extend prior work by addressing challenges (b) and (c) while also simplifying the construction of OBBs. We present DOBB, a BVH that leverages discrete rotations and shared OBB transforms at the hierarchy node level to improve traversal performance in ray tracing. Our approach is optimized for hardware-accelerated ray-OBB intersection tests. As with prior methods, we employ a post-processing step to construct DOBB-BVHs from existing AABB hierarchies, enabling seamless integration into existing BVH formats and build pipelines. The main contributions of this work are:

- The introduction of a predefined set of discrete rotations using hemisphere and angle quantization.
- A novel OBB representation that maintains a common OBB transformation shared between all wide internal node children.
- An efficient orientation mapping technique using k -DOP proxies interpreted as apex point maps [LK15].
- The usage of a compact encoding scheme that enables lossless reconstruction of OBB transformations.
- A comprehensive comparison between wide-node BVHs with shared orientations and per-child OBBs, evaluating BVH quality, traversal efficiency, and real-time performance on existing ray tracing hardware.

The remainder of this paper is structured as follows. Sec. 2, is concerned with related work. Sec. 3 introduces the creation of discrete rotations, shared rotations, and our BVH build algorithm and traversal. Sec. 4 details the implementation of our algorithm. In Sec. 5, we evaluate traversal and frame-time performance. Sec. 6 concludes this paper and discusses future work.

2. Related Work

Bounding volume hierarchies Due to the latest advances in GPU technology and hardware-accelerated ray tracing, efficient paral-

lel construction of BVHs on the GPU has been extensively surveyed [MOB*21] in recent years. Targeting GPU acceleration, linear bounding volume hierarchies [LGS*09, Kar12, Ape14] are the most common approaches to rapidly construct BVHs. Gu et al. [GHFB13] first presented an agglomerative clustering technique for fast bottom-up BVH construction. Karras and Aila [KA13] extended this agglomerative approach and proposed to restructure small treelets. Domingues and Pedrini [DP15] enhanced this approach by using a greedy clustering metric and widening binary BVHs while improving construction time with agglomerative treelet restructuring (ATRBVH). Building on clustering, Meister and Bittner [MB18a] proposed a bottom-up clustering algorithm to construct BVHs while making use of spatial proximity of neighboring primitives, called parallel locally-ordered clustering (PLOC), and introduced a parallel topology reorganization algorithm [MB18b]. Benthin et al. [BDTD22] published PLOC++, a revised version of PLOC and with a highly improved construction time. More recently, Benthin et al. [BMB*24] also introduced H-PLOC, partitioning and clustering scene geometry guided by LBBVH construction [Ape14] and invoking PLOC++ for nodes enclosing clusters with a target primitive count. They showed that it competes with ATRBVH and PLOC++ and offers good BVH quality at fast BVH construction times. Our paper uses H-PLOC to construct AABB BVHs, used as input to our OBB conversion step.

Orientation Computation Computing the optimal orientation or frame-of-reference for a given set of points or triangles has been the subject of past research since 1985. The very first algorithm to compute optimal tight-fitting bounding boxes was presented by O'Rourke [O'R85]. A more popular approach is using principal component analysis (PCA) [SGM96, KHM*98], where the predominant axes representing the largest variation in a point cloud set is determined in $\mathbb{R}^3$ space. A drawback of PCA is that it introduces bias in the resulting orientation and thus fails to capture the orientation of enclosed geometry effectively in some edge cases [DKKR09]. Larsson and Linus [LK11] introduced the ditetrahedron OBB algorithm (DiTO) to tightly fit bounds in constant time. They analyzed a set of unique triangles, constructed a base triangle from extremal points, and constructed two tetrahedra facing in opposite directions. They have demonstrated that DiTO outperforms PCA in terms of robust OBB construction and can construct nearly optimal OBB rotations for arbitrary oriented geometry.

OBB hierarchies The use of OBB hierarchies in ray tracing has gained increasing attention over the last decade. Woop et al. [WBW*14] explored the effectiveness of OBBs for hair and fur rendering by leveraging the predominant orientation of long, thin primitives for intersection tests. They constructed a hybrid BVH using both AABBs and OBBs in a top-down manner within a CPU-based ray tracer. Wald et al. [WMZ*20] utilized hardware-accelerated ray-AABB intersection tests and instance transformations to optimize ray tracing for elongated geometry. Their approach applied OBB-aware transformations at primitive nodes, demonstrating improved performance using modern APIs and GPUs. More recently, Vitsas et al. [VEPG23] introduced a method using k -DOPs as proxies for sets of points per node, guiding the construction of DiTO to establish dominant directions for optimal rotations. Their algorithm converted AABBs into OBBs as a post-processing step on the GPU in a bottom-up manner, yield-




Figure 2: Overview of our OBB post-processing algorithm: 1) For each leaf node, the best rotation is derived from the triangle or quad frame and k -DOP proxies are computed. 2) At interior nodes, the best OBB rotation candidate is selected based on the child node with the largest surface area. 3) Geometry is refit using the chosen rotation and proxies, and OBB data is stored if it improves over the AAB heuristic. 4) Proxies are combined and the OBB candidate is propagated up. Steps 2–4 repeat until a stopping criterion or the root node is reached.

ing significant ray tracing speedups for certain types of geometry. Sabino et al. [SVCNM23] constructed OBBs from k -DOPs with orthogonal bases (O-DOPs). Building on this work, Kacerik and Bittner [KB24] extended the use of k -DOP proxies by integrating them into surface-area-aware BVH construction using PLOC. They demonstrated that these k -DOPs could further improve OBB bounds, enhancing traversal efficiency and intersection performance. More recently, Kacerik and Bittner introduced skewed OBBs and custom slab tests to improve ray traversal performance. [KB25] The work in this paper is based on the idea of using k -DOP proxies and bottom-up OBB conversion. In contrast to previous research, however, our approach determines a consistent orientation for all children in an internal hierarchy node and uses much less memory to store this information while providing a SIMD-friendly ray-OBB intersection test. Furthermore, our work is inspired by the DiTO approach since we face a similar challenge of finding optimal OBBs for arbitrarily oriented geometry. As pointed out by Vitsas et al., the computational cost of exact construction algorithms is prohibitive for parallelization, while principal component analysis (PCA) suffers from limited accuracy. In contrast to their approach, we further reduce computational complexity and memory lookups by restricting OBBs to a limited set of rotations.

3. Methodology

To enable the fast generation of our DOBB-BVH, we implemented a post-processing step in a bottom-up manner following the construction of an initial AAB BVH-N, where interior nodes contain up to N children and leaf nodes enclose up to 2 triangles forming a quad. Sec. 4 covers the details of the wide AAB BVH creation.

A general overview of our designed GPU-driven post-processing approach is illustrated in Fig. 2. First, our algorithm computes the local frames of reference, leveraging topological information to efficiently select the optimal rotation from a discrete set of candidates. Since our BVH leaf nodes are restricted to two triangles, we can rapidly compute face normals and obtain shared or non-shared edges to form an initial frame. The core challenge of our work is then to identify best OBB candidates for N child nodes for all internal hierarchy nodes. We begin by precomputing k -DOPs for each leaf node in parallel. These proxies project the underlying geometry onto a limited, predefined set of axes and are stored as auxiliary sideband data during BVH construction, alongside the best-fitting OBB candidate. This information is then propagated up the hierarchy to the parent interior nodes. At each internal node,

we use a heuristic-driven strategy to select consistent OBB rotation candidates from a fixed, discrete rotation set. The OBB construction involves mapping proxies to the optimal orientation based on the propagated best-fitting OBB candidates and refitting the corresponding bounding volume extents using the underlying OBB rotation. Once computed, we proceed with the final OBB encoding if the computed OBB bounding extents offer a measurable improvement in its cost heuristic. For encoding, a compact representation using only a few bits is employed, alongside the AABs defined in the rotated OBB space. During bottom-up traversal, proxies of hierarchy nodes are updated at a higher level by merging k -DOP proxies and forwarding the best OBB candidate information. Our algorithm then repeats with identifying promising OBB candidates identification and refitting of child AABs until it reaches the hierarchy root node or satisfies a stopping criterion, such as the maximum number of levels from the leaf node.

In the following sections, we detail the selection of discrete rotations and the steps involved in OBB construction. Finally, we describe the process of decoding OBBs and performing an efficient hardware-accelerated ray-OBB intersection.

3.1. Discrete Rotation Sets

Unlike previous approaches, we defined a constrained set of discrete rotations for OBB construction. This is motivated by the observation that geometry within lower levels of a BVH tends to share similar orientations due to spatial proximity and topological similarity, making it more efficiently represented by OBBs. Our goal is to define rotation subset that balances flexibility and computational efficiency, yielding performance gains while ensuring fast build times. To achieve this, the number of possible rotations must be both controllable and pre-computable for efficient decoding. Throughout this paper, we denote an OBB rotation as 3×3 orthonormal matrix R , with $RR^T = I$. Our final set of discrete rotations $\hat{R}_D$ comprises D distinct rotations. We constructed this set by quantizing rotation angles through uniform subdivision of the 2D circle, allowing control over the number of rotations around arbitrary axes. Specifically, for a quantization step size of $\Delta = \frac{\pi}{2m}$ with up to m angles in a single direction, our angle set is defined as:

$$A := \left\{ \pm \Delta, \pm 2\Delta, \dots, \pm m\Delta \right\} \equiv \left\{ \pm \frac{\pi}{2m}, \pm \frac{\pi}{m}, \dots, \pm \frac{\pi}{2} \right\}.$$

For each axis, we selected angles in both positive and negative directions within the range $(0, \frac{\pi}{2}]$, excluding 0 to prevent identity ro-

tations. Empirically, we found that restricting rotations to $\pm \frac{\pi}{2}$ reduces redundant rotations around arbitrary axes compared to using the broader range $(0, \pi]$.

Next, we defined a rotation representation that incorporates discrete angles, generating orthogonal, non-axis-aligned frames to construct our final set $\hat{R}_D$. One common choice is the Euler angle representation, which parametrizes rotations using three quantized angles: yaw α , pitch β , and roll γ . However, the Euler angle representation has several drawbacks, including: (a) **Gimbal lock**, where one rotation angle can effectively nullify the effect of the other two. (b) **Redundant rotations**, as different combinations of Euler angles can yield identical transformations. (c) **Higher computational and numerical costs**, since constructing the final rotation requires composing three matrices $R = R_z(\gamma)R_y(\beta)R_x(\alpha)$. An alternative approach is to define an initial rotation axis using polar coordinates and apply quantized angles for the rotation around it. This method requires storing both the selected quantized angle and the rotation axis, from which the final rotation can be reconstructed as either a quaternion or a 3×3 matrix. While this approach avoids some of the pitfalls of Euler angles, it introduces additional complexity. Specifically, the flexibility in choosing the initial axis makes it more difficult to control the total number of representable rotations and requires extra steps to reconstruct the final rotation matrix.

To overcome these challenges, we predefined a fixed set of rotation axes K along with A . This strategy offers two key advantages: (a) It provides an approximately uniform distribution of axes across the 3D hemisphere, ensuring comprehensive coverage. (b) It allows for efficient precomputation of a lookup table containing $D = |K| \cdot |A|$ discrete rotations, enabling fast retrieval and reducing computational overhead. One possible way to create the set K is to use a discrete representation of a sphere, such as Platonic solids. Their well-defined vertex and face information allows us to create different combinations of unique axes along the topology. In combination with the angles in A , we eventually created table of discrete rotations. In our work, we chose up to 13 axes and use Euclidean axes along with 6 diagonal axes lying within two planes and 4 axes in 3D space. Notably, this method does not avoid duplicates in R_D as some rotations inevitably produce similar sets of orthogonal vectors. An example of how different configurations of $\hat{R}_D$ map to BVH quality is shown in Table 1 and discussed in Sec. 5.

3.2. Per-Node Geometry Projections

Our algorithm projects enclosed primitives onto k -DOP axes at each leaf node to compute initial proxy bounds. This projection step is inherently parallel, as the leaf nodes reference independent primitives and can thus be processed in isolation. For single-triangle leaves, we construct a local basis matrix $B = (\vec{a}_0, \vec{a}_1, \vec{a}_2)$, where $\vec{a}_0$ and $\vec{a}_1$ correspond to the second-longest and shortest triangle edge, and $\vec{a}_2 = \vec{a}_0 \times \vec{a}_1$, as depicted in Fig. 2. Matrix B is converted to an azimuth-angle representation and mapped to the closest orientation $R \in \hat{R}_D$ using a precomputed lookup table. For quad leaves, we derive B from the two non-shared edges, whereas for disconnected triangles, we default to the first triangle only. The initial proxy data is propagated bottom-up through the hierarchy. At each internal node, projection extents are aggregated along the same k -DOP axes by merging the minimum and maximum values per axis from all children. Parent nodes use this aggregated data during the

OBB construction to compute the child OBB extents for a candidate rotation to eventually determine whether OBBs should be applied. The projection values of our k -DOP proxies are stored as sideband data along each internal node throughout the BVH construction.

3.3. Consistent Orientation Determination

A core step in our approach involves selecting the optimal rotation $R \in \hat{R}_D$ at each internal node to minimize the cost heuristic over its child bounding boxes. Our algorithm begins by determining R for leaf nodes and propagates the selected rotation upward through the BVH tree to rapidly determine candidate rotations. While brute-force evaluation of all rotations yields optimal approach, it is impractical for real-time applications due to the high cost of recomputing bounds for all candidates and possible rotations. An alternative is the DiTO algorithm, which constructs a base triangle, identifies extremal points, and tests multiple edges. However, our experiments revealed that this method, while effective at hierarchy nodes higher up the tree, remains computationally expensive and does not consistently yield superior results compared to our approach. Instead, we found empirically that using the propagated OBB candidate from the child with the largest surface area produces acceptable results at significantly lower cost. This reduces candidate selection to a single rotation and requires only one reprojection of child node extents. With this, wave intrinsics can be used to efficiently identify the largest-area child. If the chosen R reduces the cost heuristic, we construct an OBB at the internal node and replace the AABB bounds with rotated equivalents for all children.

3.4. Bounds Recomputation




Figure 3: Schematic of OBB extent computation. The rotation basis vector $\vec{a}_i$ traverses a triangle facet of a tessellated quad-sphere, whose vertices v_i represent a k -DOP axis. Planes are constructed using those axes and the projection extent. Their intersection yields the apex point, which is projected onto $\vec{a}_i$ to obtain the final extents.

Once the candidate OBB rotation R is determined, we compute the bounding planes of all interior node children in rotated space. Contrary to [VEPG23], which recomputes the axial projection extents for k -DOPs at each internal node using all primitive vertices within the subtree, we avoid this costly operation due to its high memory access overhead. Instead, we restrict vertex projections to leaf nodes only and utilize k -DOP proxies to determine the OBB extents of interior node children. However, since the OBB basis vectors induced by R are generally misaligned with the k -DOP axes, a direct use of our proxies is not possible. This misalignment requires a conservative computation of the three bounding planes. To address this, we employ apex point maps, which enable fast and conservative fitting of bounding planes for arbitrarily oriented geometries. The technique operates on a tessellated surface of triangle

facets, in our work a triangulated quad-sphere, where the facet vertices define directions corresponding to the proxy axes. The three OBB extents for each child node are computed using this representation, as illustrated in Fig. 3. The algorithm starts by converting R to its three orthonormal basis vectors $\vec{a}_0$, $\vec{a}_1$, and $\vec{a}_2$. These vectors are tested against the quad-sphere. Once the intersecting triangle facet with vertices v_0 , v_1 , and v_2 is identified for each $\vec{a}_i$, the apex points are computed using geometric intersection tests. Specifically, we define planes using the k -DOP axes corresponding to v_0 and v_1 as normals and the projection extent as distance and perform a plane-plane intersection, followed by a plane-line intersection with the plane associated with v_2 to determine the final apex point. Final OBB extents are computed from the three apex points by projecting those onto $\vec{a}_i$ via a dot product.

3.5. OBB Cost Heuristic

Since bounding boxes are no longer constrained to align with the Euclidean axes, a key challenge lies in defining an effective cost heuristic to guide the creation of OBB nodes. We focused on the surface area heuristic (SAH), widely used in BVH research and designed to guide BVH construction such that the computed cost approximates the expected ray tracing performance, see [AKL13]. To keep our algorithm computationally efficient while maintaining effectiveness, we require a heuristic that minimizes SAH and is both lightweight and impactful. Potential options include: (a) surface area minimization, (b) child-to-child overlap reduction, and (c) approximate k -DOP SAH computation [KB24]. Approach (c) involves sampling the BVH in a top-down manner, explicitly computing the SAH for k -DOPs to assess whether an OBB reduces intersection likelihood. However, this method is computationally expensive, requires evaluation at every node and offers limited parallelism. Additionally, their approach is tailored to 14-DOPs, whereas our approach uses 26-DOPs, enabling tighter OBB construction. Approach (b) requires $(N-1)!$ computations for a node with N children and does not necessarily improve traversal performance unless normalized and combined with surface area minimization, which introduces an empirically determined weighting factor. In our experiments, minimizing the total surface area of all child bounding boxes proved to be the most effective and computationally efficient strategy, consistent with the assumption that smaller bounds lower intersection probability for arbitrary rays.

Defining an internal node N with n child nodes C_i and the surface area SA_R of their bounding boxes defined in the transformed space R , we finally determine the node cost:

$$C_R(N) = \sum_{i=0}^n SA_R(C_i). \quad (1)$$

We construct an OBB with rotation R whenever $C_R(N) < \alpha C_I(N)$, where I is the reference identity matrix and AABB frame, and α is a user-defined threshold. Although we recommend choosing α from the range $(0.9, 1]$, larger values > 1 can be used to soften the OBB creation criterion. In our experiments, we found that this could also improve performance depending on the scene and cast rays.

3.6. OBB Encoding Scheme

Keeping memory usage as low as possible is essential for real-time applications, necessitating a compact BVH node layout. Given that

all child node AABBs are defined in a common rotated space, we encode R within the internal node data using only a few additional bits. To facilitate both efficient reconstruction and minimal memory overhead, we store all rotations in $\hat{R}_D$ using two precomputed lookup tables. The first table stores a pool of shared floating-point components reused across all 3×3 rotation matrices. The second table specifies each rotation $R \in \hat{R}_D$ by indexing into the first, effectively encoding each matrix using 12 references to shared values. Since our OBB construction is limited to D rotations, each internal node only requires at most $\lceil (\log_2(D)) \rceil$ additional bits. During BVH traversal, decoding is efficient: the stored index is used to retrieve the corresponding matrix via the second table, with final floating-point values reconstructed by referencing the first table.

3.7. BVH Traversal

Algorithm 1 Our SIMD GPU kernel to traverse DOBB-BVHs.

```

 $S \leftarrow \text{Stack}([\text{root}])$ 
ray  $\leftarrow$  origin, direction
 $H \leftarrow \text{invalidNode}, \text{ray.tMax}$ 
while not  $S.\text{IsEmpty}()$  do
    node  $\leftarrow \text{FetchNode}(S.\text{PopNextHitCandidate}())$ 
    if  $\text{IsInteriorNode}(node)$  then
         $R \leftarrow \text{FetchTransform}(node.\text{obbIndex})$ 
        localRay  $\leftarrow \text{RayTransform}(\text{ray}, R)$ 
        hits  $\leftarrow \text{IntersectRayNodeChildren}(\text{localRay}, node)$ 
         $S.\text{SortAndPushCandidates}(\text{hits})$ 
    else if  $\text{IsTriangleNode}(node)$  then
        hits  $\leftarrow \text{IntersectRayTriangles}(node)$ 
         $H.\text{Update}(\text{hits})$ 
    end if
end while

```

Our BVH traversal is based on a SIMD stack-based traversal loop executed on the GPU. At each iteration, we decode the node's AABBs, child offsets, and, if present, the OBB transform. The ray is then tested for intersection against all child nodes. The traversal stack maintains the next hit candidates, ordered by a heuristic that prioritizes the objects most likely to be hit. For OBB-based BVHs, the local rotation matrix R is obtained as previously described. The input ray is transformed into the rotated space prior to intersection with all internal node children. Although reusing the transformed ray across nodes with identical rotations is theoretically possible, changes in R require applying the inverse of the previous transformation, potentially introducing numerical errors through repeated transformations. To mitigate this, our implementation consistently recomputes the local ray directly from the original one, at the cost of using additional vector registers to store the local ray. A key advantage of this approach is compatibility with existing hardware-accelerated ray-box intersection intrinsics, given that the BVH is stored in the appropriate hardware memory layout. Additionally, we leverage hardware-accelerated ray-matrix transformations to enable ray-OBB intersection. Regarding memory latency, fetching a node index requires a single memory access to load the node from VRAM or cache, while retrieving R involves an additional lookup in constant global memory.

4. Implementation

To generate wide AABB-BVHs, we first build our BVH-2 using H-PLOC [BMB*24] offering a balance between fast build times and quality. Previous studies have shown that SAH-optimal widening of BVH-2 significantly reduces traversal iterations by enabling more effective splits [AL09], improving performance on GPUs due to its SIMD-friendly nature. Therefore, a widening step is applied in a top-down manner to expand the BVH-2 nodes to our target width [YKL17] of 8 children, while optimizing for SAH. Since our leaf nodes may contain up to two triangles, we aim to form quads. Quadification can be performed either as a pre-processing step by forming triangle pairs while minimizing the surface area of their bounding box, or during the BVH widening phase by merging two leaf nodes enclosed by a single interior node.

Based on our studies, we opted for $\hat{R}_{104}$ as the discrete rotations set for our final performance studies with $|K| = 13$ axes and $|A| = 8$ angles ($m = 4$). This allows us to encode an OBB index with just 7 bits at each interior node. The algorithm is implemented as a single HLSL kernel based on shader model 6.3 in Direct3D12 that executes after the construction of a BVH-N and functions as a bottom up post-process pass. It starts with an initial list of internal nodes that have only leaves as children, and use N threads per element in this list, where N is the number of children per internal node. Every internal node is then processed with up to N threads to pick the child with largest surface area and its best $R \in \hat{R}_{104}$. To track and recompute child node bounds, we use 26-DOPs stored along with each node. After refitting of bounds and computing the cost heuristic, our kernel encodes the final OBB index and bounds if Eq. 1 is met. Refitting fits within a work-group processor and local cache while utilizing a single fetch instruction to speed up execution. In our traversal kernel, we make use of a single HW intrinsic performing both ray transform and intersection based on the encoded OBB index and bounding boxes for an N wide interior hierarchy node.

In addition to our GPU-based approach, we developed a CPU-based reference version of our algorithm to estimate the theoretical upper bound of potential performance improvements achievable with DOBB-BVHs. Unlike the GPU variant, this version performs a brute-force evaluation of all possible OBB candidates. For each candidate rotation, it computes projection extents along all unique 225 basis vectors derived from the full set of rotations in $\hat{R}_{104}$. This approach eliminates the need of apex point maps and enables exact computation of bounding box extents in the rotated space. Rather than selecting the OBB candidate from the child node with the largest surface area, the CPU algorithm evaluates the complete set of rotations, computes the corresponding bounding volumes, and selects R that minimizes the cost function.

All our kernels and following performance measurements were run on an AMD Radeon RX 9070 XT GPU with 16GB DDR6 video memory and boost frequency up to 2.9 GHz. CPU BVH builds were executed on an AMD Ryzen 9 5950X processor.

5. Analysis

We evaluated our approach in terms of traversal behavior, memory consumption, build time, and real-time performance. To compare our GPU algorithm to speed-of-light BVH builds and OBB construction, we make use of our offline CPU BVH post-processing

		Total Number of Rotation Angles and Discretization Steps (m)							
		2(1)	4(2)	6(3)	8(4)	10(5)	12(6)	14(7)	16(8)
Total Number of Rotation Axes	3	195.6	167.2	158.0	144.6	141.8	149.5	148.3	147.5
	4	195.6	167.2	141.9	132.6	140.8	140.0	141.5	147.5
	5	173.9	148.8	133.3	127.3	130.3	137.6	138.6	137.7
	6	173.9	148.8	131.2	124.5	124.9	128.7	135.0	137.7
	7	173.9	141.6	128.8	122.6	122.6	124.5	127.6	133.0
	8	173.9	141.6	125.9	120.6	120.8	122.2	124.6	126.9
	9	163.6	136.5	123.4	119.2	119.7	120.7	122.0	124.7
	10	163.6	136.5	123.2	118.4	118.2	119.0	120.2	121.6
	11	163.6	132.2	122.7	117.5	116.9	117.9	118.9	119.8
	12	163.6	132.2	121.8	116.7	116.2	116.9	117.6	119.0
	13	153.5	128.9	121.3	116.0	115.6	115.7	116.5	117.6

Table 1: SAH computed on DOBB-based BVH-8 generated from Hairball with OBB_{soGPU} using different axes and angles counts, highlighting axis-aligned orientations only and promising configs.

step as explained in Sec. 4, applied to the pregenerated AABB BVH-Ns. We further use a branching factor $N \in 4, 6, 8$ in our studies to demonstrate the benefit of our algorithm for wide BVHs. In the remainder of this section, we refer to OBB_{PC} with the construction of OBBs for each single child encoded in sideband data alongside all hierarchy nodes. OBB_{soCPU} denotes the CPU-based method, while OBB_{soGPU} represents our GPU-based approach.

5.1. Datasets

For our studies, we selected 5 scenes with different levels of complexity, as shown in Table 3. *Hairball* is a scene with hair-like geometry and randomly oriented hair strands well suited for OBBs and posing a challenge for AABB BVHs to well represent the underlying geometry. The *White Oak* is a tree with triangulated leaves of different orientation in 3D space. *Dragon* represents a mixture of axis and non-axis aligned geometry where parts of the model, such as the sword and tail, will benefit from OBBs particularly well. *Bistro Exterior* scene is comprised of the same amount of triangles as hairball but with a mixture of axis-aligned and rotated buildings, furniture, and foliage. *Powerplant* contains massive triangle geometry and some non-axis aligned geometry, such as construction crossbars that can lead to highly increased iteration counts due to multi-intersection caused by multiple ill-fitting AABBs.

5.2. BVH Quality

We first evaluated the quality of BVHs by measuring SAH. Prior research has shown that SAH approximately correlates with the final ray tracing performance for rays cast from the outside of the geometry [AKL13]. For both OBB approaches, we compute an approximate SAH by taking the child AABB stored in the parent node as base to compute the SAH of its subtree. To determine an appropriate configuration for $\hat{R}_D$, we evaluated various discrete orientation sets R_D with $|A| \in [3, 13]$ and $m \in [1, 8]$, where $|K| = 2m$, with consistently oriented hair strands. We measured the resulting SAH values of DOBB-BVHs constructed for the *Hairball* scene using a branching factor of $N = 8$ (see Table 1). As expected, the SAH reduction improves with an increased number of discrete rotations. However, the improvements diminish beyond 10 unique axes. Notably, a configuration with $m = 1$ and all three Euclidean axes can be neglected as a 90-degree rotation simply inverts the axes.

Based on this analysis, we selected $D = 13 \cdot 8 = 104$ discrete rotations as a balanced trade-off between construction quality and memory overhead. Using $\hat{R}_{104}$, we then evaluated SAH values across all five datasets and multiple BVH branching factors (refer to Table 3). The most significant SAH reductions occur in the *Hairball* and *Bistro Exterior* scenes, where $\text{OBB}_{\text{soCPU}}$ achieves a 40–47% SAH reduction compared to AABB BVH-Ns, consistently across all branching factors. In comparison, $\text{OBB}_{\text{soGPU}}$ reduces SAH by approximately 20–25%. The *Powerplant* dataset benefits similarly from oriented bounding volumes. $\text{OBB}_{\text{soGPU}}$ achieves a 15–17% SAH reduction, only about 10% below the results of the CPU-based offline builder. Conversely, the *White Oak* and *Dragon* scenes reveal limitations of shared OBBs. These scenes feature characteristics such as small primitives, highly axis-aligned geometry, or randomized orientations, which makes it challenging to identify fitting OBBs. In these cases, SAH reduction is limited to 14%, with most results in the 2–6% range. This trend worsens with increasing branching factors, highlighting the challenge of coherent orientation fitting in wide BVHs. Overall, $\text{OBB}_{\text{soCPU}}$ delivers performance comparable to per-child orientation approach used in OBB_{PC} , particularly in scenes with strong orientation coherence.

5.3. Traversal Behavior

We analyzed the effect of BVH quality on traversal efficiency by measuring per-ray iteration steps needed to find the closest triangle hit. On our target GPU, which supports HW-accelerated ray intersection with 8 OBB child nodes and 2 triangles per step, iteration count corresponds directly to node visits. Since long-running rays can stall entire SIMD wavefronts, we evaluated both maximum and average iteration counts (see Table 3) and visualized traversal cost per pixel for BVH-8 with primary and global illumination (GI) rays (Tables 4 and 8, App. B). The most notable gains were observed in maximum iteration counts. For *Hairball*, *Powerplant*, and *Dragon*, $\text{OBB}_{\text{soCPU}}$ reduced max. iterations by 40–80%, and $\text{OBB}_{\text{soGPU}}$ by 30–40%. In *Dragon*, even minor SAH improvements helped addressing poorly approximated tail regions. Similarly, in *Powerplant*, rays traversing non-axis-aligned crossbars caused over 1000 iterations, significantly mitigated by OBBs. GI rays showed even greater improvements, with max. iteration counts up to $3 \times$ higher than primary rays. $\text{OBB}_{\text{soCPU}}$ reduced these by up to 91% for *Powerplant* and 82% for *Hairball* (Table 7). We found that average iteration reductions are less pronounced. For primary rays, $\text{OBB}_{\text{soCPU}}$ achieved 20–25% improvements, and $\text{OBB}_{\text{soGPU}}$ 17–22%. Axis-aligned scenes like *White Oak* saw limited gains (1–8%), as orientation fitting is less effective. This gets increasingly challenging for wider BVHs, where a consistent orientation becomes harder to identify. Therefore, OBB_{PC} generally outperforms shared-orientation DOBBs. Still, $\text{OBB}_{\text{soCPU}}$ delivers comparable results, improving average traversal by 10–20% over $\text{OBB}_{\text{soGPU}}$.

5.4. Memory Consumption

We store k floats for each plane of a given k -DOP, defaulting to 26 floats for a total of 104 bytes. Our worst-case allocation is $(N - 1) \cdot 104b$ for $N - 1$ interior nodes. We also maintained a list of sorted node indices from which to start the refit kernel, which requires $N \cdot 4$ bytes where N is the maximum number of internal nodes. Finally, 3 lookup tables were used for our precomputed set of orientations.

Scene	Approach	Build Time ms	Primary GRays/s	Secondary GRays/s
Hairball	AABB	11.62	2.91	1.16
	$\text{OBB}_{\text{soGPU}}$	12.82 (1.10×)	4.16 (1.43×)	1.66 (1.43×)
	$\text{OBB}_{\text{soCPU}}$	-	5.10 (1.76×)	2.16 (1.86×)
Bistro Ext.	AABB	12.16	3.63	1.34
	$\text{OBB}_{\text{soGPU}}$	13.56 (1.12×)	4.27 (1.18×)	1.97 (1.47×)
	$\text{OBB}_{\text{soCPU}}$	-	5.01 (1.38×)	2.77 (2.06×)
White Oak	AABB	0.65	7.95	2.82
	$\text{OBB}_{\text{soGPU}}$	0.71 (1.09×)	8.04 (1.01×)	2.93 (1.04×)
	$\text{OBB}_{\text{soCPU}}$	-	8.23 (1.04×)	3.06 (1.08×)
Dragon	AABB	2.57	20.93	3.26
	$\text{OBB}_{\text{soGPU}}$	3.02 (1.18×)	21.08 (1.01×)	3.37 (1.03×)
	$\text{OBB}_{\text{soCPU}}$	-	20.87 (1.00×)	3.47 (1.06×)
Powerplant	AABB	51.02	4.09	1.87
	$\text{OBB}_{\text{soGPU}}$	55.20 (1.08×)	5.32 (1.30×)	3.09 (1.65×)
	$\text{OBB}_{\text{soCPU}}$	-	4.96 (1.21×)	3.68 (1.96×)

Table 2: Performance of our CPU- and GPU-based DOBB BVH-8 build algorithm compared against an AABB tree and the relative factor within brackets (increased build time or improved performance).

To map our OBB matrix index to a 3×3 rotation, we relied on a 2-level lookup table that uses 960B total, and to map the rotation to an OBB matrix index, we use a 3D look up table that uses 4096B.

5.5. Build Time and Real-Time Performance

We evaluated the build time of our $\text{OBB}_{\text{soGPU}}$ approach including all AABB build passes and measured final ray tracing performance by tracing both primary and secondary rays through each scene from 25 different viewpoints. Results for BVH-8 are presented in Table 2 and highlighted in Fig. 1. Overall, $\text{OBB}_{\text{soGPU}}$ introduces a moderate build time overhead of approximately 10–15% relative to the baseline BVH construction pipeline. In absolute terms, this corresponds to an additional 0.1 ms for the smallest scene and up to 5 ms for the largest scene. We consider this acceptable for real-time applications, particularly for static content. In return, our DOBB-BVHs deliver notable run-time improvements. For primary rays, we observe an average speedup of 18.5%, with gains reaching up to 43% for primary rays. The most substantial improvements appear for secondary rays, particularly in global illumination (GI) workloads where rays originate from geometry surfaces and exhibit high divergence. In these cases, performance increases by up to 65%. In comparison, $\text{OBB}_{\text{soCPU}}$ achieves promising traversal speedups of up to 76% for primary rays and 106% for GI rays. These results reinforce our earlier findings: scenes with predominantly non-axis-aligned geometry benefit most, due to significant reductions in maximum traversal iterations. As supported by both SAH and iteration count analyses, these performance gains hold consistently across all tested BVH branching factors and ray types. A breakdown of BVH construction phases is provided in Fig. 4, while Table 6 in App. A contains details build times and rendering performance across all branching factors.

6. Conclusion

In this paper, we have introduced DOBB-BVH, a novel bounding volume hierarchy based on OBBs, which maintains the same rotation between all wide hierarchy node children through a fixed set





Scene	Approach	Surface Area Heuristic (SAH)			Max. Iterations / Ray			Avg. Iterations / Ray		
		BVH-4	BVH-6	BVH-8	BVH-4	BVH-6	BVH-8	BVH-4	BVH-6	BVH-8
	AABB (Baseline)	254.6	216.7	195.6	403	340	304	49.3	42.0	37.7
	AABB → OBB _{PC}	141.5	115.3	101.7	226	187	158	34.0	28.8	25.8
	(Roofline Per-Child)	0.56×	0.53×	0.52×	0.59×	0.63×	0.64×	0.73×	0.73×	0.74×
	AABB → OBB _{soCPU}	153.5	128.3	116.0	236	213	196	35.9	30.6	27.7
	Ours (Brute-Force)	0.60×	0.59×	0.59×	0.59×	0.63×	0.64×	0.73×	0.73×	0.74×
	AABB → OBB _{soGPU}	200.7	171.2	146.7	314	265	215	41.8	36.0	31.4
	Ours (GPU-based)	0.79×	0.79×	0.75×	0.78×	0.78×	0.71×	0.85×	0.86×	0.83×
	AABB (Baseline)	48.5	42.5	39.7	237	200	175	43.0	35.8	31.8
	AABB → OBB _{PC}	29.9	22.3	21.3	194	195	184	34.7	29.5	26.1
	Rel. Factor to AABB	0.62×	0.52×	0.54×	0.82×	0.98×	1.05×	0.81×	0.82×	0.82×
	AABB → OBB _{soCPU}	26.7	22.7	20.8	207	183	151	38.1	31.7	28.1
		0.55×	0.53×	0.52×	0.87×	0.92×	0.86×	0.89×	0.89×	0.88×
	AABB → OBB _{soGPU}	38.7	32.9	29.6	230	199	179	40.7	34.5	30.8
		0.80×	0.77×	0.75×	0.97×	1.00×	1.02×	0.95×	0.96×	0.97×
	AABB (Baseline)	67.9	58.3	53.8	156	137	119	20.0	17.5	16.3
	AABB → OBB _{PC}	57.0	48.4	46.7	117	112	94	17.6	15.6	14.4
		0.84×	0.83×	0.87×	0.75×	0.82×	0.79×	0.88×	0.89×	0.88×
	AABB → OBB _{soCPU}	61.9	53.4	49.5	134	118	110	18.9	16.8	15.7
		0.91×	0.92×	0.92×	0.86×	0.86×	0.92×	0.95×	0.96×	0.97×
	AABB → OBB _{soGPU}	66.1	56.8	52.6	150	132	113	19.7	17.3	16.1
		0.97×	0.98×	0.98×	0.96×	0.96×	0.95×	0.98×	0.99×	0.99×
	AABB (Baseline)	25.9	21.7	19.7	208	176	153	12.3	10.5	9.8
	AABB → OBB _{PC}	21.4	17.7	13.0	90	75	68	10.9	9.4	8.6
		0.82×	0.81×	0.66×	0.43×	0.43×	0.44×	0.89×	0.89×	0.88×
	AABB → OBB _{soCPU}	22.4	18.8	17.0	94	81	71	11.3	9.8	9.0
		0.86×	0.86×	0.86×	0.45×	0.46×	0.46×	0.92×	0.93×	0.92×
	AABB → OBB _{soGPU}	24.6	20.6	18.6	161	124	107	12.0	10.2	9.4
		0.95×	0.95×	0.95×	0.77×	0.70×	0.70×	0.97×	0.97×	0.97×
	AABB (Baseline)	20.1	17.5	16.3	1788	1606	1574	48.1	38.5	36.6
	AABB → OBB _{PC}	14.3	12.2	11.2	369	326	306	36.2	28.3	26.6
		0.71×	0.69×	0.69×	0.21×	0.20×	0.19×	0.75×	0.73×	0.73×
	AABB → OBB _{soCPU}	14.6	12.5	11.5	373	333	315	38.1	29.4	27.5
		0.72×	0.71×	0.71×	0.21×	0.21×	0.20×	0.79×	0.76×	0.75×
	AABB → OBB _{soGPU}	17.1	14.5	13.5	831	725	565	39.5	30.6	28.5
		0.85×	0.83×	0.83×	0.46×	0.45×	0.36×	0.82×	0.79×	0.78×

Table 3: BVH quality and traversal statistics for primary rays across different BVH widths for all datasets and OBB conversion algorithms.

of discrete rotations. Our approach employs an efficient GPU-based post-processing algorithm to identify and compute OBBs from existing AABB BVHs, leveraging hardware-accelerated ray-OBB intersections. To evaluate the effectiveness of DOBB-BVH, we have conducted a detailed performance study demonstrating significant improvements in ray traversal: up to an 80% reduction in worst-case traversal steps and a 10–20% decrease in average iterations per ray. Despite extra a build-time overhead of 10–12%, our method achieves real-time ray tracing performance improvements of 22.7% on average, with gains up to 43% for primary rays and 65% for secondary rays. These results are enabled by an efficient mapping of geometry to optimal OBB candidates and an improved refitting strategy using apex point maps as proxies.

In the future, our primary objective is to close the noticeable performance gap of ~20% in traversal efficiency and ~40% in overall performance compared to a brute-force CPU baseline. Additional promising research directions include developing improved cost heuristics, optimizing the mapping of geometry and child nodes to OBB candidates, and refining the discrete rotation set. We additionally plan to combine the implementation of DOBB with

triangle splitting. We also see potential in hybrid bounding volume approaches, such as incorporating AABBs and OBBs within a single node (e.g., 12-DOPs), to better adapt to heterogeneous scene content. Modifying early stages of the BVH build process can yield benefits. For example, constructing OBB-aware quads or orientation-coherent clusters could improve fitting quality. Finally, we aim to investigate efficient OBB construction techniques tailored for dynamic or updatable BVHs, enabling broader applicability in real-time and interactive rendering scenarios.

Acknowledgements We thank our colleague Julien Beasley for his work on look-up tables. The *Dragon* model was obtained from Benedikt Bitterli's Rendering Resources site, the *Bistro Exterior* model is courtesy of Amazon Lumberyard. All other scenes were obtained from Morgan McGuire's resource archive. Datasets and bounding box images were rendered using the AMD Capsaicin Framework [BOM*23] and Mitsuba [Wen10].

Copyright Notice and Trademarks

© 2025 Advanced Micro Devices, Inc. All rights reserved. AMD, Ryzen, Radeon, and combinations thereof are trademarks of Ad-




Table 4: Visualization of traversal loop iteration counts at each pixel for primary rays for our CPU-based and GPU-based DOBB BVHs compared to the baseline AABB BVH.

vanced Micro Devices, Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.

References

- [AKL13] AILA T., KARRAS T., LAINE S.: On quality metrics of bounding volume hierarchies. In *Proceedings of the 5th High-Performance Graphics Conference* (New York, NY, USA, 2013), HPG '13, Association for Computing Machinery, p. 101–107. [5](#), [6](#)
- [AL09] AILA T., LAINE S.: Understanding the efficiency of ray traversal on gpus. In *Proceedings of the Conference on High Performance Graphics 2009* (New York, NY, USA, 2009), HPG '09, Association for Computing Machinery, p. 145–149. [6](#)
- [AMD25] AMD: RDNA4 Instruction Set Architecture, 2025. <https://amd.com/content/dam/amd/en/documents/radeon-tech-docs/instruction-set-architectures/rdna4-instruction-set-architecture.pdf> [Accessed: 2025-05-19]. [1](#)
- [Ape14] APETREI C.: Fast and simple agglomerative LBVH construction. In *Computer Graphics and Visual Computing* (2014), Borgo R., Tang W., (Eds.), The Eurographics Association. [2](#)
- [BDTD22] BENTHIN C., DRABINSKI R., TESSARI L., DITTEBRANDT A.: PLOC++: Parallel locally-ordered clustering for bounding volume hierarchy construction revisited. *Proc. ACM Comput. Graph. Interact. Tech.* 5, 3 (7 2022). [2](#)
- [BMB*24] BENTHIN C., MEISTER D., BARCZAK J., MEHALWAL R., TSAKOK J., KENSLE A.: H-PLOC: Hierarchical parallel locally-ordered clustering for bounding volume hierarchy construction. *Proc. ACM Comput. Graph. Interact. Tech.* 7, 3 (Aug. 2024). [2](#), [6](#)
- [BOM*23] BOISSÉ G., OLIVER M., MEUNIER S., DUPONT DE DINECHIN H., ETO K.: *The AMD Capsaicin Framework*. 8 2023. https://github.com/Radeon-Pro/Capsaicin_Open [Accessed: 2025-05-19]. [8](#)
- [BKRR09] DIMITROV D., KNAUER C., KRIEGLER K., RÖTE G.: Bounds on the quality of the PCA bounding boxes. *Computational Geometry* 42, 8 (2009), 772–789. [2](#)
- [DP15] DOMINGUES L. R., PEDRINI H.: Bounding volume hierarchy optimization through agglomerative treelet restructuring. In *Proceedings of the 7th Conference on High-Performance Graphics* (New York, NY, USA, 2015), HPG '15, Association for Computing Machinery, p. 13–20. [2](#)
- [Eri04] ERICSON C.: *Real-Time Collision Detection*. The Morgan Kaufmann Series in Computer Graphics. Taylor & Francis, 2004. [2](#)
- [GHFB13] GU Y., HE Y., FATAHALIAN K., BLELLOCH G.: Efficient BVH construction via approximate agglomerative clustering. In *Proceedings of the 5th High-Performance Graphics Conference* (New York, NY, USA, 2013), HPG '13, Association for Computing Machinery, p. 81–88. [2](#)
- [INT24] INTEL: Intel Arc graphics developer guide for real-time ray tracing in games, 2024. <https://cdrdv2-public.intel.com/766679/intel-rttr-applications-developer-guide-v4.pdf> [Accessed: 2025-05-26]. [1](#)
- [KA13] KARRAS T., AILA T.: Fast parallel construction of high-quality bounding volume hierarchies. In *Proceedings of the 5th High-Performance Graphics Conference* (New York, NY, USA, 2013), HPG '13, Association for Computing Machinery, p. 89–99. [2](#)
- [Kar12] KARRAS T.: Maximizing parallelism in the construction of BVHs, octrees, and k-d trees. In *Proceedings of the Fourth ACM SIGGRAPH / Eurographics Conference on High-Performance Graphics* (Goslar, DEU, 2012), EGGH-HPG'12, Eurographics Association, p. 33–37. [2](#)
- [KB24] KÁČERIK M., BITTNER J.: SAH-optimized k-DOP hierarchies for ray tracing. *Proc. ACM Comput. Graph. Interact. Tech.* 7, 3 (Aug. 2024). [2](#), [3](#), [5](#)
- [KB25] KÁČERIK M., BITTNER J.: SOBB: Skewed oriented bounding boxes for ray tracing. *Computer Graphics Forum* (2025). [3](#)
- [KHM*98] KLOSOWSKI J., HELD M., MITCHELL J., SOWIZRAL H., ZIKAN K.: Efficient collision detection using bounding volume hierarchies of k-dops. *IEEE Transactions on Visualization and Computer Graphics* 4, 1 (1998), 21–36. [2](#)
- [LGS*09] LAUTERBACH C., GARLAND M., SENGUPTA S., LUEBKE D., MANOCHA D.: Fast bvh construction on gpus. *Computer Graphics Forum* 28, 2 (2009), 375–384. [2](#)
- [LK11] LARSSON T., KÄLLBERG L.: *Fast Computation of Tight-Fitting Oriented Bounding Boxes*. CRC Press, 02 2011, pp. 3–19. [2](#)
- [LK15] LAINE S., KARRAS T.: Apex point map for constant-time bounding plane approximation. In *Eurographics Symposium on Rendering - Experimental Ideas and Implementations* (2015), Lehtinen J., Nowrouzezahrai D., (Eds.), The Eurographics Association. [2](#)
- [MB18a] MEISTER D., BITTNER J.: Parallel locally-ordered clustering for bounding volume hierarchy construction. *IEEE Transactions on Visualization and Computer Graphics* 24, 3 (2018), 1345–1353. [2](#)
- [MB18b] MEISTER D., BITTNER J.: Parallel reinsertion for bounding volume hierarchy optimization. *Computer Graphics Forum* (2018). [2](#)
- [MOB*21] MEISTER D., OGAKI S., BENTHIN C., DOYLE M., GUTHE M., BITTNER J.: A survey on bounding volume hierarchies for ray tracing. *Computer Graphics Forum* 40 (06 2021), 683–712. [1](#), [2](#)
- [NVI25] NVIDIA: NVIDIA RTX Blackwell GPU Architecture, 2025. <https://images.nvidia.com/aem-dam/Solutions/geforce/blackwell/nvidia-rtx-blackwell-gpu-architecture.pdf> [Accessed: 2025-05-19]. [1](#)
- [O'R85] O'ROURKE J.: Finding minimal enclosing boxes. *International Journal of Computer & Information Sciences* 14 (1985), 183–199. [2](#)
- [RW80] RUBIN S. M., WHITTED T.: A 3-dimensional representation for fast rendering of complex scenes. In *Proceedings of the 7th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1980), SIGGRAPH '80, Association for Computing Machinery, p. 110–116. [2](#)
- [SE02] SCHNEIDER P., EBERLY D. H.: *Geometric Tools for Computer Graphics*. The Morgan Kaufmann Series in Computer Graphics. Elsevier Science, 2002. [2](#)
- [SGM96] S. GOTTSCHALK M. C. L., MANOCHA D.: OBBTree: a hierarchical structure for rapid interference detection. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1996), SIGGRAPH '96, Association for Computing Machinery, p. 171–180. [2](#)
- [SVCNM23] SABINO R., VIDAL C. A., CAVALCANTE-NETO J. B., MAIA J. G. R.: Building oriented bounding boxes by the intermediate use of ODOPs. *Computers and Graphics* 116, C (2023), 251–261. [3](#)
- [VEPG23] VITSAS N., EVANGELOU I., PAPAIOANNOU G., GKARAVELIS A.: Parallel transformation of bounding volume hierarchies into oriented bounding box trees. *Computer Graphics Forum* (2023). [2](#), [4](#)
- [WBW*14] WOOP S., BENTHIN C., WALD I., JOHNSON G. S., TABELLION E.: Exploiting local orientation similarity for efficient ray traversal of hair and fur. In *Proceedings of High Performance Graphics* (Goslar, DEU, 2014), HPG '14, Eurographics Association, p. 41–49. [2](#)
- [Wen10] WENZEL J.: Mitsuba renderer, 2010. <http://www.mitsuba-renderer.org>. [8](#)
- [WMZ*20] WALD I., MORRICAL N., ZELLMANN S., MA L., USHER W., HUANG T., PASCUCI V.: Using hardware ray transforms to accelerate ray/primitive intersections for long, thin primitive types. *Proc. ACM Comput. Graph. Interact. Tech.* 3, 2 (aug 2020). [2](#)
- [YKL17] YLITIE H., KARRAS T., LAINE S.: Efficient incoherent ray traversal on GPUs through compressed wide BVHs. In *Proceedings of High Performance Graphics* (New York, NY, USA, 2017), HPG '17, Association for Computing Machinery. [6](#)

Appendix A: Breakdown of BVH Build Phases

We measured the execution times of all individual steps in our BVH construction pipeline, as shown in Fig. 4, and report the results in milliseconds in Table 5. For scenes containing up to three million triangles, our DOBB post-processing step takes up to 1.15ms and increases to approximately 5ms for scenes with 13 million triangles.




Figure 4: Breakdown of BVH-8 build phases across our set of test scenes, showcasing the impact of DOBB-BVH.

Build Phases	White Oak	Dragon	Bistro Ext.	Hairball	Powerplant
Encode Primitives	0.39 ms	0.33 ms	1.04 ms	1.36 ms	4.62 ms
Morton Codes	0.01 ms	0.05 ms	0.26 ms	0.58 ms	2.05 ms
Merge Sort	0.04 ms	0.13 ms	1.21 ms	1.38 ms	7.29 ms
H-PLOC	0.16 ms	0.51 ms	1.67 ms	2.67 ms	5.94 ms
Encode HW Format	0.09 ms	0.86 ms	2.77 ms	3.07 ms	11.32 ms
DOBB Pass	0.05 ms	0.35 ms	1.13 ms	1.15 ms	4.95 ms

Table 5: Results of BVH-8 build phases in Figure 4.

Appendix B: Analysis of Traversal Performance for Primary and Secondary Rays

We further evaluated the build times and traversal performance of our DOBB algorithm for both primary and secondary rays. Table 6 presents the build and traversal performance across BVHs with varying branching factors. Detailed analysis of BVH quality and traversal behavior for secondary (global illumination) rays is provided in Table 7 and Table 8.

Scene	Approach	Build Time ms			Primary GRays/s			Secondary GRays/s		
		BVH-4	BVH-6	BVH-8	BVH-4	BVH-6	BVH-8	BVH-4	BVH-6	BVH-8
Hairball	AABB	12.61	12.07	11.62	2.27	2.63	2.91	0.89	1.06	1.16
	OBB _{soGPU}	14.96 (1.19×)	13.26 (1.10×)	12.82 (1.10×)	3.19 (1.40×)	3.62 (1.38×)	4.16 (1.43×)	1.43 (1.35×)	1.66 (1.43×)	1.66 (1.43×)
Bistro Ext.	AABB	13.49	12.31	12.16	2.82	3.29	3.63	1.14	1.29	1.34
	OBB _{soGPU}	15.57 (1.26×)	13.56 (1.10×)	13.56 (1.12×)	3.22 (1.14×)	3.78 (1.15×)	4.27 (1.18×)	1.56 (1.36×)	1.82 (1.41×)	1.97 (1.47×)
White Oak	AABB	0.70	0.55	0.65	6.57	7.49	7.95	2.39	2.70	2.82
	OBB _{soGPU}	0.76 (1.09×)	0.687 (1.25×)	0.71 (1.09×)	6.64 (1.01×)	7.61 (1.02×)	8.04 (1.01×)	2.48 (1.04×)	2.80 (1.03×)	2.93 (1.04×)
Dragon	AABB	2.57	2.49	2.57	17.78	19.99	20.93	2.87	3.17	3.26
	OBB _{soGPU}	3.22 (1.25×)	3.09 (1.24×)	3.02x (1.18×)	18.19 (1.02×)	20.41 (1.02×)	21.08 (1.01×)	3.00 (1.04×)	3.26 (1.03×)	3.37 (1.03×)
Powerplant	AABB	51.02	49.50	48.46	3.54	3.92	4.09	1.73	2.11	1.87
	OBB _{soGPU}	60.33 (1.18×)	55.37 (1.12×)	55.20 (1.08×)	5.04 (1.42×)	6.37 (1.63×)	5.32 (1.30×)	2.29 (1.33×)	3.08 (1.46×)	3.09 (1.65×)

Table 6: Hardware results of our GPU-based DOBB BVHs build algorithm across all branching factors 4–8 compared against an AABB tree and the relative factor within brackets (*increased build time* or *improved performance*).





	Scene [#Triangles]	Approach	Max. Iterations / Ray			Avg. Iterations / Ray		
			BVH-4	BVH-6	BVH-8	BVH-4	BVH-6	BVH-8
	Hairball [2.9M]	AABB (Baseline)	3735	409	378	34.3	28.9	25.9
		AABB → OBB _{PC}	212	159	151	20.8	17.2	15.4
		(Roofline Per-Child)	0.06×	0.39×	0.40×	0.61×	0.60×	0.60×
		AABB → OBB _{soCPU}	346	180	176	22.2	18.7	16.9
		Ours (Brute-Force)	0.09×	0.44×	0.47×	0.65×	0.65×	0.65×
		AABB → OBB _{soGPU}	1050	268	247	27.4	23.4	20.1
		Ours (GPU-based)	0.28×	0.66×	0.65×	0.80×	0.81×	0.78×
	Bistro Exterior [2.8M]	AABB (Baseline)	524	492	463	61.4	52.4	48.3
		AABB → OBB _{PC}	253	220	209	39.3	32.9	29.6
		Rel. Factor to AABB	0.48×	0.45×	0.45×	0.64×	0.63×	0.61×
		AABB → OBB _{soCPU}	256	225	213	40.8	34.2	30.7
			0.49×	0.46×	0.46×	0.66×	0.65×	0.64×
		AABB → OBB _{soGPU}	361	317	316	52.5	43.5	39.3
			0.69×	0.64×	0.68×	0.86×	0.83×	0.81×
	White Oak [36K]	AABB (Baseline)	189	157	147	15.2	13.0	12.0
		AABB → OBB _{PC}	155	111	104	12.7	10.9	10.0
			0.82×	0.71×	0.71×	0.84×	0.84×	0.84×
		AABB → OBB _{soCPU}	170	131	121	13.8	12.0	11.2
			0.90×	0.83×	0.82×	0.91×	0.93×	0.93×
		AABB → OBB _{soGPU}	184	146	128	14.7	12.6	11.6
			0.97×	0.93×	0.87×	0.97×	0.97×	0.97×
	Dragon [830K]	AABB (Baseline)	226	195	194	9.0	7.5	6.8
		AABB → OBB _{PC}	120	98	91	7.6	6.4	5.8
			0.53×	0.50×	0.47×	0.85×	0.85×	0.84×
		AABB → OBB _{soCPU}	109	90	81	7.8	6.6	6.0
			0.48×	0.46×	0.42×	0.87×	0.88×	0.88×
		AABB → OBB _{soGPU}	154	131	110	8.6	7.1	6.5
			0.68×	0.67×	0.57×	0.96×	0.95×	0.95×
	Powerplant [12.8M]	AABB (Baseline)	2892	2625	2061	31.6	26.1	23.2
		AABB → OBB _{PC}	447	467	475	25.8	20.9	18.4
			0.15×	0.18×	0.23×	0.82×	0.80×	0.79×
		AABB → OBB _{soCPU}	510	471	425	26.8	21.6	18.9
			0.18×	0.18×	0.21×	0.85×	0.83×	0.82×
		AABB → OBB _{soGPU}	1573	1358	971	27.4	22.3	19.6
			0.54×	0.52×	0.47×	0.87×	0.85×	0.84×

Table 7: BVH quality and traversal statistics for secondary (GI) rays across different BVH widths for all datasets and OBB conversion algorithms.




Table 8: Visualizaton of traversal loop iteration counts at each pixel for secondary (GI) rays for our CPU-based and GPU-based DOBB BVHs compared to the baseline AABB BVH.