

Exploiting Ray Tracing Technology Through OptiX to Compute Particle Interactions with Cutoff in a 3D Environment on GPU

David Algis¹, Bérenger Bramas²
University of Poitiers, XLIM, France¹
Studio Nyx, France¹
Inria Nancy, France²
ICube Laboratory, France²
University of Strasbourg, France²

Abstract—Particle interaction simulation is a fundamental method of scientific computing that require high-performance solutions. In this context, computing on graphics processing units (GPUs) has become standard due to the significant performance gains over conventional CPUs. However, since GPUs were originally designed for 3D rendering, they still retain several features that are not fully exploited in scientific computing. One such feature is ray tracing, a powerful technique for rendering 3D scenes. In this paper, we propose exploiting ray tracing technology via OptiX and CUDA to compute particle interactions with a cutoff distance in a 3D environment on GPUs. To this end, we describe algorithmic techniques and geometric patterns for efficiently determining the interaction lists for each particle. Our approach enables the computation of interactions with quasi-linear complexity in the number of particles, eliminating the need to construct a grid of cells or an explicit kd-tree. We compare the performance of our method to a classical grid-based approach and demonstrate that our approach is faster in most cases with non-uniform particle distributions.

Keywords—CUDA; graphics processing unit; high-performance computing; OptiX; particle interactions; ray tracing; scientific computing

I. INTRODUCTION

High-performance computing (HPC) is a key technology in scientific computing. Since the early 2000s, the use of graphics processing units (GPUs) has become standard in HPC, and they now equip many of the fastest supercomputers.¹ GPUs are massively parallel processors that allow computations to be performed on thousands of cores, making them perfectly suited for inherently parallel problems. The use of GPUs in scientific computing has led to incredible performance gains in many fields, such as molecular dynamics [1], fluid dynamics [2], astrophysics [3], and machine learning [4].

Most scientific computing applications that use GPUs are based on the CUDA² programming model, and to a lesser extent, the OpenCL programming model. They do not exploit all the features of GPUs, particularly those dedicated to 3D rendering. Among these features, ray tracing is a powerful technology used to render 3D scenes. In this method, rays are cast from the camera to the scene, and the intersections

of the rays with the objects in the scene are computed to generate the colors of the pixels in the image [5]. For example, the NVIDIA GeForce RTX 4090, as a single consumer-grade GPU, demonstrates the raw power of this technology by rendering complex 3D scenes at 87 frames per second in 4K resolution (3840 x 2160 pixels), handling millions of triangles per frame [6].

In this paper, we are interested in evaluating how ray tracing technology could be used to compute particle interactions in a 3D environment. Our motivation is twofold: We want to evaluate if it is possible to use ray tracing technology, and we want to create the algorithmic patterns needed for that purpose.

With these aims, we focus on the computation of particle interactions in a 3D environment, which is a common problem in scientific computing. When the potential of the interaction kernel decreases exponentially with distance, the interactions can be computed with a cutoff distance, i.e. the interactions are only computed between particles that are closer than a given distance, achieving less but still satisfactory accuracy. This allows the complexity of the interactions to drop from $O(N^2)$ to $O(N)$, where N is the number of particles if the cutoff distance is small enough. Classical methods to compute such interactions are based on the use of a grid of cells or an explicit kd-tree to quickly find the interaction lists for each particle. In this paper, we aim to avoid using such data structures and instead exploit ray tracing technology.

The contributions of this paper are as follows:

- We propose a method to compute particle interactions with a cutoff distance in a 3D environment based on ray tracing technology.
- We describe two algorithmic techniques based on geometric patterns to find the interaction lists for each particle using real intersections.
- We compare the performance of our approach with a classical approach based on a grid of cells and with an existing method that relies on ray tracing.

The rest of the paper is organized as follows: In Section II, we present the prerequisites. In Section III, we review the related work. In Section IV-A, we present our proposed

¹<https://top500.org/>

²<https://developer.nvidia.com/cuda-toolkit>

approach. In Section V, we present the performance study. Finally, we conclude in Section VI.

II. PREREQUISITES

A. Particle Interactions

Computing the interactions between N particles in a 3D environment is a common problem in scientific computing. For example, this is essential in fluid simulations using smoothed particle hydrodynamics [7] and in molecular dynamics for simulating forces between atoms [8]. These interactions are usually modeled by a potential function that depends on the distance between the particles. A straightforward way to compute the interactions is to evaluate the potential function for all pairs of particles. However, the potential function can be short-range or long-range. When the potential function is short-range, the interactions can be computed with a cutoff distance, i.e. the interactions are only computed between particles that are closer than a given distance. This reduces the complexity of the interactions from $O(N^2)$ to $O(N)$, where N is the number of particles, but this is only possible if we have an efficient way to find the particles' neighbors. Moreover, the positions of the particles are usually updated after each computation step. Consequently, the system used to find the interactions between the particles should be updated at each iteration of the simulation.

A possible solution to get the interaction list is to build a grid of cells mapped over the simulation space, where each cell contains the particles that are inside. The cells have a width equal to the cutoff distance C . Then, for each particle, the interactions are computed with the particles that are in the same cell and in the neighboring cells. Building the grid of cells and finding the interaction lists for each particle can be done in $O(N)$: we start by computing the cell index for each particle, then we order the particles in a new array by assigning a unique index per particle using atomic operations, and finally, we move the particles to a new array [9]. Each of these three operations can be implemented with a parallel loop over the N particles.

B. Graphics Processing Units

Graphics processing units (GPUs) are massively parallel processors that allow computations to be performed on thousands of cores. The hardware design of GPUs has been optimized for graphics rendering, particularly for the rendering of 3D scenes. To this end, GPUs have features dedicated to 3D rendering, such as texture mapping, rasterization, and ray tracing. Internally, GPUs are organized in a hierarchy of processing units, including streaming multiprocessors (SMs), warp schedulers, and execution units.

NVIDIA has proposed the CUDA programming model to develop parallel applications for GPUs. CUDA is designed to express algorithms in a way that can be mapped to the GPUs' hardware organization. For instance, thread blocks are distributed across SMs, and threads are executed in warps. There are also keywords and functions to use shared memory, constant memory, and texture memory. Thus, creating optimized applications for GPUs requires an understanding of the hardware architecture of GPUs and potentially adapting algorithms to their specificities.

C. Ray Tracing

Among the many features of GPUs dedicated to 3D rendering, ray tracing stands out as a powerful technology widely used in video games for rendering realistic 3D scenes [10]. It is a hardware-accelerated technique that computes the interactions of rays with objects in a scene. Ray tracing generates pixel colors in an image by casting rays from the camera into the scene. For each intersection of a ray with an object, the pixel's color is determined based on the object's material properties and lighting conditions. Rays can also be reflected or refracted by objects, or continue through non-opaque surfaces, enabling recursive computation of interactions for enhanced realism.

Typically, ray tracing kernels are implemented within the shaders of the graphics pipeline. Shaders are small programs executed on the GPU for each pixel in an image, often written in specialized languages such as GLSL (for OpenGL) or HLSL (for DirectX). These shaders run in parallel on the GPU, allowing for concurrent computation of ray-object interactions. NVIDIA introduced a method to utilize ray tracing within the CUDA programming model through its OptiX library [11]. With OptiX, developers retain the flexibility of CUDA programming while leveraging ray tracing technology, albeit with some constraints on kernel implementation.

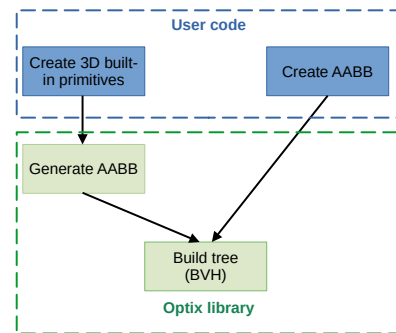


Fig. 1. Schematic view of BVH tree creation. The OptiX library requires a list of encompassing axis-aligned bounding boxes (AABBs), which can either be directly provided or generated from a list of basic primitives.

In ray tracing, scenes are represented using geometric primitives such as triangles or spheres, which are encapsulated within bounding volumes to enhance computational efficiency [12]. A common bounding volume is the Axis-Aligned Bounding Box (AABB), a rectangular box aligned with the coordinate axes and defined by its minimum and maximum bounds $(x_{min}, y_{min}, z_{min})$ and $(x_{max}, y_{max}, z_{max})$. AABB-ray intersection tests are conducted by calculating t_{min} and t_{max} for each axis and ensuring that overlaps exist across all axes, efficiently determining whether the ray intersects the box.

To further optimize performance, AABBs are organized into a Bounding Volume Hierarchy (BVH), a tree-like data structure where each node represents an AABB. Internal nodes group child AABBs, while leaf nodes encapsulate the actual primitives. In OptiX, users can either provide a list of basic primitives or directly supply a list of AABBs, as illustrated in Fig. 1.

During ray traversal, the algorithm performs intersection tests at each BVH node. If a ray misses an AABB, the entire

subtree beneath that node is skipped, reducing unnecessary computations. When a ray reaches a leaf node, precise intersection tests are conducted with the enclosed primitives, and the closest intersection is recorded. This process reduces the complexity of ray-scene intersections from linear to logarithmic by pruning large portions of the scene, ensuring only relevant branches of the BVH are explored.

The final output of the ray tracing algorithm, typically the closest hit point, is used for shading or further processing to determine the visual appearance of the scene. This hierarchical approach, combined with the use of AABBs and BVH, ensures that ray tracing remains computationally feasible even for complex scenes containing millions of primitives.

Fig. 2 provides a summary of the different operations in ray tracing using OptiX. In the first step, the user casts rays in the desired direction, and OptiX manages the BVH tree traversal and potential interactions. When a ray enters an enclosing bounding box (AABB), a callback function is invoked to determine whether the ray actually intersects the corresponding primitive, based on a built-in or user-defined Intersection Shader (IS), referred to as the intersection strategy. For custom primitives, the user must implement the IS. If the ray does not intersect or is marked to continue, the process is repeated for other primitives. The traversal stops when a final intersection is found or no more primitives can be intersected, triggering a call to the miss callback.

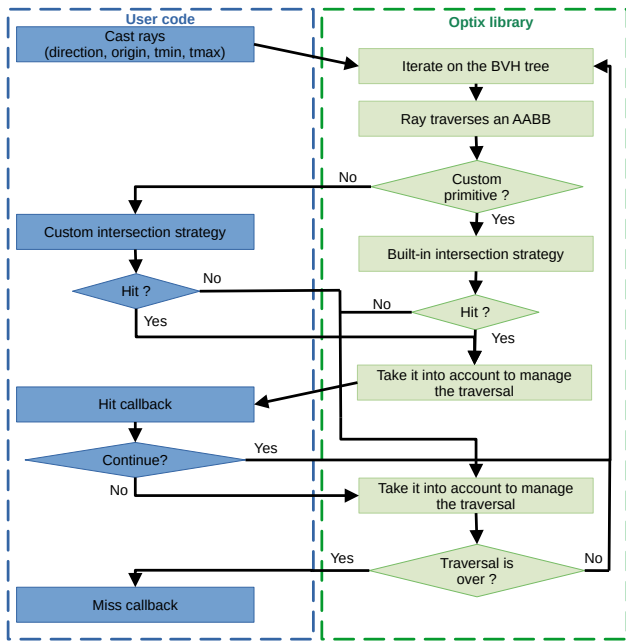


Fig. 2. Schematic view of the OptiX ray tracing workflow.

III. RELATED WORK

A. Neighbor Search on GPU

The main work on physical simulation of particle interactions on GPU has been proposed by [13], to compute the gravitational potential. They described an efficient implementation using shared memory that became a standard implementation on GPU. As mentioned in Section II-A, grid-based and kd-tree structures are widely utilized for neighbor

search operations on GPUs. For grid-based approaches, a detailed description can be found in our previous study [9], which focuses on scenarios with few particles per cell and demonstrates that utilizing shared memory often does not yield significant benefits. Regarding kd-tree implementations, [14] present a method for fast k-nearest neighbor searches using GPUs, highlighting the efficiency of kd-trees in high-dimensional spaces.

B. Ray Tracing in Computer Graphics

Ray tracing has been widely applied in particle-based representations, primarily in the domain of computer graphics and rendering. The author in [15] propose a method for efficiently ray tracing Gaussian particles to enable advanced rendering effects such as shadows, reflections, and depth of field in dense particle scenes, with applications in novel-view synthesis and visual realism. Similarly, [16] explores hardware-accelerated ray tracing for rendering particles that cast shadows, focusing on evaluating the performance of a prototype system. While these works demonstrate the utility of ray tracing in particle-based scenes, their focus lies in rendering and visualization, contrasting with our application of ray tracing for neighbor search in physical simulations.

C. OptiX in Scientific Computing

The NVIDIA OptiX framework has shown potential for diverse applications in scientific computing. Blyth et al. utilized OptiX to enable high-performance optical photon simulations in particle physics. This approach reduced memory and computation overheads by using GPU-based culling of photon hits, handling millions of photons in complex geometries [17], [18]. OptiX has been utilized as a flexible and high-performing tool for optical 3D modeling, enabling virtual measurements of sample surfaces by tracing over 1 billion light rays per image and comparing simulated results with those from physical measuring devices [19].

While these applications highlight OptiX's potential in handling computationally intensive tasks, they primarily focus on modeling physical processes rather than using ray tracing for spatial queries or particle interactions. Our work bridges this gap by leveraging OptiX for neighbor detection and particle interaction calculations, building on its demonstrated strengths in scientific modeling to expand its applicability into the domain of particle-based simulations.

D. OptiX for Neighbor Search

Using OptiX to find neighbors between elements has already been explored in several works. I. Evangelou et al. [20] introduced a novel approach to spatial queries, particularly radius search, by leveraging GPU-accelerated ray-tracing frameworks with OptiX. Instead of traditional spatial data structures like kd-trees, the authors proposed mapping the radius-search task to the ray-tracing paradigm by treating query points as primitives within a bounding volume hierarchy (BVH). In the proposed method, a ray used in the radius-search operation is essentially infinitesimal in extent, and its purpose is not to compute a traditional intersection but to test whether the origin of the ray (the query point) is within the bounding volume of a "sphere" surrounding each sample point. For that purpose, the

authors reimplemented the function that tests if a ray intersects a sphere to control how the tests are performed. This approach enabled significant performance gains in dynamically updated datasets.

Yuhao Zhu [21] further advanced the application of ray tracing for neighbor search by introducing optimizations for mapping the neighbor search problem onto the ray-tracing hardware available in modern GPUs. The author identified two key performance bottlenecks: unmanaged query-to-ray mapping, which led to control-flow divergences, and excessive tree traversals stemming from monolithic BVH construction. To address these issues, they proposed query scheduling and partitioning strategies that exploit spatial coherence and reduce BVH traversal time. Their experiments demonstrated substantial speedups ranging from 2.2 to 65.0 over existing GPU neighbor search libraries. Their work highlights the potential of using ray-tracing hardware not only for rendering but also for efficient spatial queries. The source code for their work is available on GitHub; however, the implementation is not directly applicable to our purpose as it primarily focuses on radius search and constructing interaction lists. In contrast, our interest lies in directly computing the interactions. Furthermore, for large test cases involving particles distributed on a sphere, the performance of their method appears to be of the same order of magnitude as ours.

Shiwei Zhao et al. [22] extended the use of ray tracing for particle-based simulations by converting neighbor search into a ray tracing problem. Each particle was represented as a bounding box, similar to previous works, with tiny rays emitted to detect intersections and identify neighboring particles. By leveraging NVIDIA's RT cores alongside CUDA cores, they demonstrated 10% to 60% performance improvements over traditional cell-based methods in various particle-based simulations, including discrete element methods and smooth particle hydrodynamics. Their approach underscores the versatility of ray-tracing cores for accelerating computationally intensive neighbor search tasks across different domains.

These works collectively demonstrate the potential of using OptiX and ray-tracing hardware for efficient spatial queries, providing a foundation for further exploration in GPU-accelerated computational methods. However, these approaches lack a direct integration of physical interaction computations within the ray-tracing framework. Our approach addresses this gap by embedding cutoff distances directly into the intersection tests. Additionally, we introduce two novel methods for neighbor detection that utilize actual intersection computations, accommodating scenarios where custom IS are unsupported.

IV. PROPOSED SOLUTIONS

A. Overview

The core idea of our approach is to represent particles using geometric primitives and to use ray tracing to find the neighbors of each particle by detecting intersections with these primitives.

With this aim, we investigate three possibilities:

1) *Custom AABB*: in this case, we directly provide the list of englobing bounding boxes. Therefore, we also have to provide our own IS in which we do not compute real intersections but only check if a ray is inside an AABB. This strategy is the closest one to the state of the art.

2) *Spheres*: in this case, we use built-in sphere primitives, which allows us to use the built-in IS that check if a ray really intersect with the surface of a sphere. However, we had to create our own geometric algorithm to make it work.

3) *Triangles*: in this case, we use built-in triangle primitives to create squares, which allows to use the built-in IS that check if a ray really intersect with the surface of a triangles. As for the spheres, we had to create our own geometric algorithm to make it work.

The spherical approach is simpler and more intuitive than using squares, but we are interested in evaluating if the squares made of triangles is more efficient as it is the most used primitive in 3D rendering. We also want to evaluate if the built-in IS is more efficient than our custom IS.

In the cases that rely on built-in IS, we use the following algorithmic pattern:

- 1) We build a geometric representation for each particle.
- 2) We cast rays from each particle in specific directions to find the neighbors, depending on the representation.
- 3) We filter the intersections to avoid computing the same interaction multiple times.
- 4) We compute the interactions between the particles that are closer than C .

To reduce overhead, we aim to use as few rays as possible and ensure they do not intersect with too many particles that are not within the distance C .

In terms of implementation, we use the OptiX library to develop the ray tracing kernels, which can be used in conjunction with the CUDA programming model. Specifically, in the OptiX API, we create a scene by providing a list of geometric primitives. We then provide a CUDA kernel that launches the rays, where each ray has an origin, a direction, a starting point, and an endpoint. Usually, one CUDA thread is used to launch one ray. Finally, a callback is invoked by OptiX when a ray intersects with a primitive or if no intersection is found between the starting and ending points (in 3D rendering, this usually means that the background color should be used).

The data accessible from the callback is limited. OptiX built-in IS can provide information about the intersection, such as the intersection point, the normal on the surface, the distance from the ray's origin, and the index of the primitive that was hit. Additionally, the user can pass information from the CUDA kernel that launches the rays to the callback using a payload. A payload is a user-defined data structure that is passed along with a ray as it traverses the scene. It allows the ray to carry information that can be read or modified. The number of payload variables is limited (usually 16 32-bit integers in recent versions).

When the hit callback is invoked, and we want the ray to continue, there are two possibilities. The first is to launch a new ray from the intersection point in a new direction. This is

done by storing the intersection distance in a payload variable, returning from the callback, and then launching a new ray from the intersection point using a loop in the CUDA kernel to reach the desired distance. The second possibility is to inform OptiX that we want to continue the traversal by calling a corresponding function in the intersection callback. This second approach is generally more efficient, as it avoids the overhead associated with launching new rays and should be favored in practice.

Additionally, we cannot allocate memory in the callbacks, so we cannot build complex data structures, such as lists, to store intersection lists. Consequently, if we want to filter the intersections, we cannot fill an array with indices and check if an index exists in the array to ensure uniqueness; instead, we must use geometric properties to filter the intersections.

In the remainder of the section, we consider that the target particle is the particle for which we want to find the neighbors.

B. Spherical Representation

In this section, we consider the case where the particles are represented by spheres and we use the built-in IS. In OptiX, a sphere is a geometric primitive defined by its center and radius, and multiple spheres of the same radius can be instantiated in the scene, which is the approach we use.

We have the following objectives:

- 1) Expressing the radius of the spheres depending on the cutoff distance;
- 2) Defining the origins and directions of the rays;
- 3) Providing a mechanism to filter the intersection when the rays intersect multiple times with the same sphere.

In our model, we will use three rays for each particle, one in each direction of the coordinate system. Consider a sphere of radius C centered at the origin in a three-dimensional space. The points on this sphere that are at the farthest distance from the three coordinate axes are located in the corners of a cube inscribed within the sphere. For instance, one such point at distance C from the origin lies in the direction $(1, 1, 1)$. These 8 points, corresponding to the vertices of the cube, are all at a distance C from the origin, and we want to know how far they are from the coordinate axes. This can be calculated as follows: Since the points have coordinates where $|x| = |y| = |z|$, we use the equation of the sphere $x^2 + y^2 + z^2 = C^2$. If we take the point for which $x = y = z$, it gives $3x^2 = C^2$, resulting in $x = \frac{C}{\sqrt{3}}$. Therefore, each of these points is at a distance of $\frac{C \times \sqrt{2}}{\sqrt{3}}$ from any of the three coordinate axes.

We use this information to define the radius of the spheres and the length of our rays. The radius is set to $r = \frac{C \times \sqrt{2}}{\sqrt{3}}$. In this scenario, it is sufficient that the rays go up to $l = \frac{C}{\sqrt{3}}$ in each direction relatively to the particle's position (so a single ray goes from $-l$ to l). We provide a simplified 3D rendering of the spheres in Fig. 3 that illustrate our model.

However, if $l < r$, there are positions where the sphere could simply englobe the rays, and we would miss some intersections (when the source and target are closer than $r - l$ in the three dimensions). Therefore, we set $l = r$ and add an ϵ to the radius of the sphere to ensure that the rays intersect with

the sphere in all cases, obtaining $r = \frac{C \times \sqrt{2}}{\sqrt{3}} + \epsilon$ and $l = \frac{C \times \sqrt{2}}{\sqrt{3}}$. The ϵ is a small value such that it must be impossible that two particles are closer than ϵ , or some intersections will be missed (see Appendix A for more details).




Fig. 3. 3D Spherical representation of a source and target particles distance from $C = 1$. The source sphere has a radius of $\frac{\sqrt{2}}{\sqrt{3}}$ and the rays, represented by segments, are of length $\frac{1}{\sqrt{3}}$ in each direction. In this case, the sphere could englobe the rays.

When the source and target particles are perfectly aligned on one axis, the ray will intersect for particles distant from the extremity by $r + \epsilon$. Therefore, two particle distant from $l + r + \epsilon$ can interact. This case and any intermediate situation where the source/target are actually too far can easily be filtered by checking the distance. In Fig. 4, we provide a 2D representation of the particles using spheres on intricate cases.

However, each ray can potentially intersect with a sphere several times, and the same sphere can be intersected by several rays of the same target particle, so we need to filter them. It is impossible to maintain a global list that all the rays can access to check if an intersection has already been found, or even a single list per ray to ensure that it does not intersect with the same sphere. Therefore, we must do this based on geometric rules as described in Algorithm 1. When we detect an intersection, we get the position of the source and target particles and first check they are within the cutoff distance. Then, we check if the closest ray to the source particle is the current one, and if yes, we can proceed with the computation (see Appendix A for more details).

C. Double Squares Representation

Most 3D rendering applications use triangles to represent the objects in the scene, which motivated us to create a second model that relies on triangles instead of spheres. Of course, building the AABB representations and the HBV tree is not expected to be faster than the sphere, but the hardware modules and built-in could be more optimize for the triangles, as it is more common primitives. In our model, we use four triangles to draw two squares, which are positioned opposite each other. Each square has a width of $C + \epsilon$ and is positioned at a distance of $C/2$ along the X-axis relative to the particle, one in each direction. We provide Code 1 in Appendix B, which shows how we generate the triangles from the particles' positions.

We then launch four rays, all with the same length and direction, but positioned at the corners of the squares. Each ray is positioned at $-C/2$ from the center of the square and has a length of $C + \epsilon$. The ϵ is used to ensure that when the source and target particles have the same x coordinate and their squares overlap, the rays cross the triangles. We provide




Fig. 4. 2D Spherical representation of the particles in three different cases. In the first one (left), what will be the position of the farthest source particle and how it will be detected by the ray. In the second one (center), we show the case where the source particle is the farthest from the rays (it also shows that in 2D the sphere radius could be smaller). In the last one (right), we show different source particles with their spheres and the rays that will intersect with them.

Algorithm 1: Sphere intersection callback

```

Data: Optix variables
Result: Callback when a ray intersect with a sphere
1 Function callback()
2 begin
   /* Get the center of the sphere (source
   position) */
3    $q \leftarrow \text{optixGetSphereData}()$ 
   /* Current particle position (target
   position) */
4    $point \leftarrow \text{getPayloadPartPos}()$ 
   /* Compute differences and distances */
5    $diff\_pos \leftarrow \{abs(point.x - q.x), abs(point.y - q.y), abs(point.z - q.z)\}$ 
6    $diff\_pos\_squared \leftarrow \{diff\_pos.x^2, diff\_pos.y^2, diff\_pos.z^2\}$ 
7    $dist\_squared \leftarrow diff\_pos\_squared.x + diff\_pos\_squared.y + diff\_pos\_squared.z$ 
   /* Get the cutoff distance from payload */
8    $c \leftarrow \text{getPayloadC}()$ 
   /* Ensure it is in the cutoff distance */
9   if  $dist\_squared < c^2$  then
10     $dist\_axis\_squared \leftarrow \{diff\_pos\_squared.y + diff\_pos\_squared.z, diff\_pos\_squared.x + diff\_pos\_squared.z, diff\_pos\_squared.x + diff\_pos\_squared.y\}$ 
11     $ray\_dir \leftarrow \text{optixGetWorldRayDirection}()$ 
12     $closest\_axis \leftarrow \text{getClosestAxis}(dist\_axis\_squared)$ 
13     $closest\_axis\_is\_ray\_dir \leftarrow (ray\_dir == closest\_axis)$ 
   /* Ensure this ray and this intersection
   are the good one */
14    if  $closest\_axis\_is\_ray\_dir$  then
       /* Call computation kernel */

```

in Fig. 5 the 2D representation of the particles using squares (which are lines in 2D).

From this description, one particle can be seen as a box of sides $(C + \epsilon, C + \epsilon, C)$ (see Fig. 6). The rays can be seen as the four edges of the box in the x direction, and the triangles composed the front and back faces. If any two boxes have an

intersection, we will detect it as shown in Fig. 5.

Potentially, the ray will intersect with the squares of the target particle, but this can easily be filtered by checking either the coordinates or the index of the geometric elements. Additionally, if the target and source particles are aligned on the y or z axis, two rays will intersect with the source's squares. To filter these intersections, we proceed as shown in Algorithm 2. We compare the coordinates between the source and the target and proceed as follows: If y and z are different, we perform the computation (only the current ray will intersect). If y is equal, we use the ray of index 0 if z is smaller, and the ray of index 2 if z is greater. If z is equal, we use the ray of index 0 if y is smaller, and the ray of index 1 if y is greater. Otherwise, only the ray of index 0 will be used for computation (all four rays will intersect).

D. Custom AABB

In this strategy, a bounding box with a width equal to $2 \times C$ is created around each particle. OptiX treats these boxes as custom primitives, requiring us to implement a custom IS that is invoked when a ray enters an AABB. We do not perform any intersection tests within the IS, as our goal is not to compute ray-AABB or ray-surface intersections but to identify pairs of particles that are within a distance C of each other. To achieve this, rays are cast from the particles' positions with an infinitesimally short length, as illustrated in Fig. 7.

The interaction between the source and target particles is computed directly within the IS callback. Consequently, the hit and miss callbacks are left empty, as there is no need to filter intersections. In the IS callback, OptiX provides information about the ray (specifically, the starting point, which corresponds to the position of the source particle) and the index of the AABB being tested. To obtain geometric information about the AABB, one could register a Primitive Geometry Acceleration Structure (PGAS) for each AABB during the OptiX scene build stage and retrieve this data in the IS callback. However, we observed that this approach significantly increased the scene build cost. Instead, we opted to access the global memory directly to retrieve the position of the target particle, which proved to be a more efficient solution.

V. PERFORMANCE STUDY

This section presents a performance evaluation of the different methods under various configurations. The analysis begins with a description of the experimental setup, including




Fig. 5. Double squares (line) representation of the particles. On the left, we show how the squares can overlap for particles that are too far, but which can be easily filtered with the distance. In the middle, we show how particles that have the same x coordinate can have their squares that overlap. On the right, we show different source particles with their squares and the rays that will intersect with them.




Fig. 6. 3D Rectangular representation of a source and target particles. Some rays (thick lines) intersect with the squares.




Fig. 7. Custom AABB representation of the particles. The rays cast from the particles' positions with an infinitesimal length. In the IS callback invoked if a ray enters an AABB, the distance is checked and the interaction between the source and target particles is computed. Self interactions are filtered by ensuring that the source and target are different.

the hardware and software used for implementation. Performance is then examined for two types of particle distributions: uniform and non-uniform. Differentiating these distributions allows for assessing the efficiency of grid-based methods like CUDA, which are optimized for uniform distributions but can encounter inefficiencies with sparse or highly localized data in non-uniform scenarios.

A. Experimental Setup

1) *Hardware*: We have used two NVidia GPUs:

- A100³ with 40GB hBM2, 48KB of shared-memory, 108 multi-processors, zero RT Cores, 8192 CUDA Cores max single-precision performance 19.5TFLOPS, and max tensor performance 311.84TFLOPS.

³<https://www.nvidia.com/content/dam/en-zz/Solutions/Data-Center/a100/pdf/nvidia-a100-datasheet-nvidia-us-2188504-web.pdf>


- RTX8000⁴ with 48GB GDDR6, 48KB of shared-memory, 72 multi-processors, 72 RT Cores, 4608 CUDA Cores, maximum Ray casting of 10Giga Rays/sec, max single-precision performance 16.3TFLOPS, and max tensor performance 130.5TFLOPS.

Despite the lack of RT cores, the A100 is capable of executing ray tracing kernels, the GPU then use its other units to behave similarly.

2) *Software*: We have implemented the proposed approach in OptiX 8.⁵ We use the GNU compiler 11.2.0 and the NVidia CUDA compiler 12.3. The source code is available online.⁶

The code was compiled with the following flags: `-arch=sm_75` for the RTX8000, `-arch=sm_80` for the A100 (and `-O3 -DNDEBUG` on both). We execute each kernel five times and take the average as reference.

We provide in Fig. 8 the 3D rendering of the primitives using ray tracing: in Fig. 8a for the spheres and in Fig. 8b for the squares. These figures were drawn using the OptiX API and ray tracing from camera to the scene.



(a) 3D Rendering of sphere primitives. (b) 3D Rendering of square primitives.

Fig. 8. 3D Rendering of the primitives using conventional ray tracing.

B. Uniform distribution

In this test case, the particles are distributed uniformly within a unit box. Consequently, there are no (or very few)

⁴<https://www.nvidia.com/content/dam/en-zz/Solutions/design-visualization/quadro-product-literature/quadro-rtx-8000-us-nvidia-946977-r1-web.pdf>

⁵<https://developer.nvidia.com/rtx/ray-tracing/optix>

⁶<https://gitlab.inria.fr/bramas/particle-interaction-with-optix>

Algorithm 2: Triangle intersection callback

```
Data: Optix variables  
Result: Callback when a ray intersect with a triangle  
1 Function callback()  
2 begin  
   /* Get information on the intersected  
   triangle */  
3 vertices ← optixGetTriangleVertexData  
   (gas, prim_idx, sbtGASIndex, 0.f, vertices)  
   /* Retrieve the source position from the  
   triangle vertices */  
4 Declare q as float3  
5 q.y ←  
   (max(vertices[0].y, vertices[1].y, vertices[2].y)+  
   min(vertices[0].y, vertices[1].y, vertices[2].y))/2  
6 q.z ←  
   (max(vertices[0].z, vertices[1].z, vertices[2].z)+  
   min(vertices[0].z, vertices[1].z, vertices[2].z))/2  
7 c ← getPayloadC()  
8 if (prim_idx mod 4) < 2 then  
9   | q.x ← vertices[0].x + c/2  
10 else  
11   | q.x ← vertices[0].x - c/2  
   /* Get target particle position */  
12 point ← getPayloadPartPos()  
   /* Compute distance */  
13 dist_p1_p2 ← distance(point, q)  
   /* Ensure it is not a self intersection and  
   it is in the cutoff distance */  
14 if dist_p1_p2 < c AND dist_p1_p2 > ε then  
15   ray_idx ← getPayloadRayidx()  
16   is_ray_for_compute ← (point.y ≠ q.y AND  
   point.z ≠ q.z)  
17   OR ((point.z < q.z AND ray_idx == 0) OR  
   (point.z > q.z AND ray_idx == 2))  
18   OR ((point.y < q.y AND ray_idx == 0) OR  
   (point.y > q.y AND ray_idx == 1))  
19   OR ray_idx == 0  
20   if is_ray_for_compute then  
   | /* Perform the computation */
```

empty cells in the grid of the CUDA version. The cutoff distance is defined as $1/\beta$, where β can be 2, 4, 8, 16, or 32. For a given β , the simulation grid consists of 3^β cells. The number of particles N is then calculated as $N = p \times 3^\beta$, where p represents the average number of particles per cell, taking values of 1, 2, 4, 8, 16, or 32. All computations are performed in single-precision floating point.

We present the results in Fig. 9. For all configurations, we measured the initialization step (light color) and the computation step (dark color).

For the OptiX-based implementation, the initialization step involves building the scene by invoking the OptiX API to create the primitives. For the CUDA version, the initialization step involves constructing the grid of cells. Consequently, in

the OptiX-based version, the computation step includes the time spent launching the rays, executing the callback functions, and performing the interactions. In the CUDA version, the computation step corresponds to the kernel time required to compute the interactions.

1) *Comparison between triangles and spheres:* First, we analyze the performance difference between triangles (blue) and spheres (green). On the A100 GPU, both models deliver similar performance, but the ratio of initialization time (light color) to computation time (dark color) is higher for spheres. This indicates that the computation step is more efficient for spheres than for triangles. This suggests that in scenarios where the initialization step is performed only once (e.g. static elements), spheres might be a better choice. However, on the RTX8000, triangles are faster than spheres across all configurations. Although the initialization step for spheres is quicker than for triangles in cases with fewer particles, this trend does not hold for larger configurations.

2) *Comparison between custom AABB and built-in primitives:* Next, we compare the Custom AABB method (orange/red) with the built-in primitives (triangles in blue and spheres in green). On the A100, the Custom AABB consistently outperforms the built-in primitives, and its advantage grows as the number of particles or cells increases. For $\beta = 32$ (Fig. 9i), the Custom AABB is up to four times faster than the triangle- or sphere-based methods. On the RTX8000, the Custom AABB performs similarly to triangles for small particle counts but becomes faster as the number of particles increases. Overall, the Custom AABB is the fastest method. This demonstrates that creating a custom Intersection Shader (IS) that is significantly simpler and lighter than the built-in ones, which compute actual intersections, can accelerate execution.

3) *Benefit of sorting the particles in the custom AABB methods:* Sorting the particles to position them closer in memory based on their spatial proximity in the simulation is expected to optimize memory accesses. However, sorting the particles also incurs a computational cost. We observe that the strategy without sorting (orange) is faster than the strategy with sorting (red) when there are few particles. However, as the number of particles increases, the benefits of sorting become significant. In our test case, for a given β , increasing the number of particles leads to an increase in neighboring particles and, consequently, the number of interactions. In such cases, optimizing memory access becomes critical.

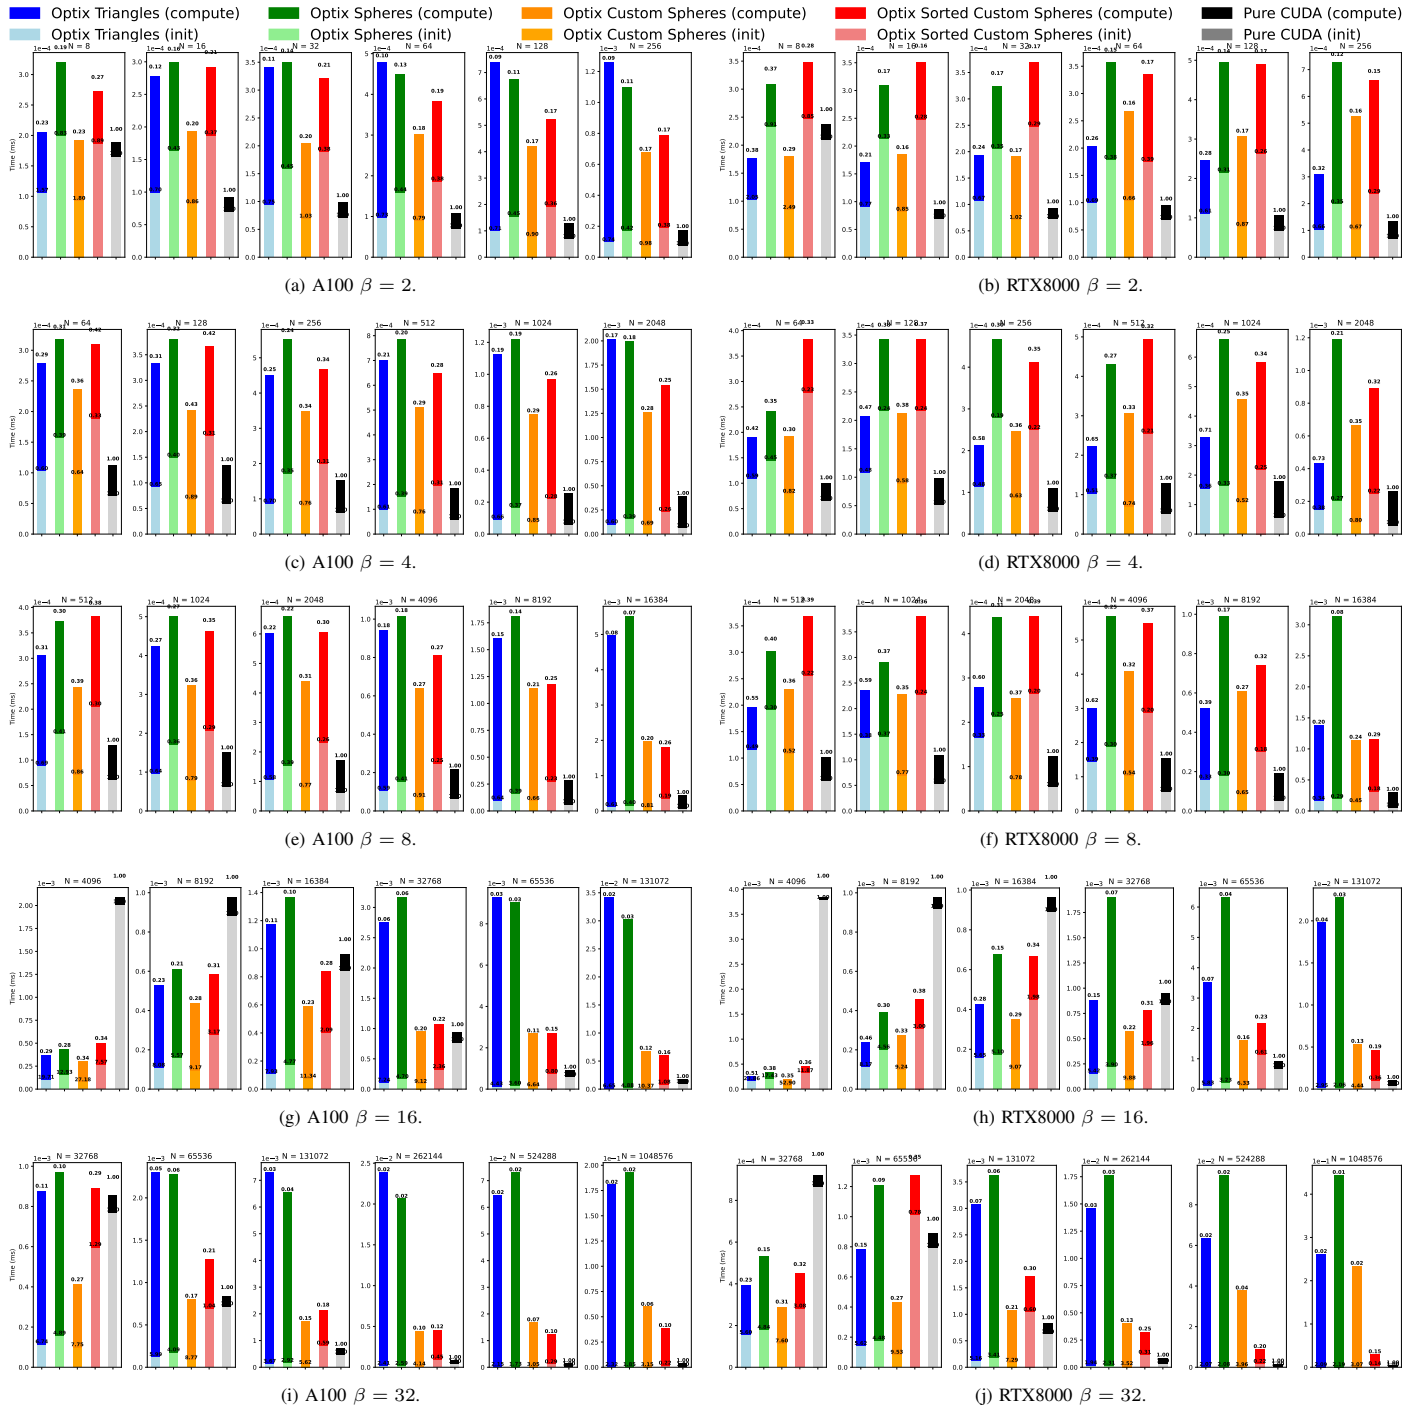


Fig. 9. Performance results for the two GPUs for our approaches (OptiX spheres, OptiX triangles, and Custom AABB without and with sort) and the pure CUDA implementation that use a grid of cells (CUDA). β is the divisor coefficient of the simulation box. The cutoff distance is $1/\beta$ and there are β^3 cells in the grid in the Cuda version. The speedup against the CUDA version is shown above the bars for both the build and compute steps.




Fig. 10. Performance results for the two GPUs for our approaches (OptiX spheres, OptiX triangles, and Custom AABB without and with sort) and the pure CUDA implementation that use a grid of cells (CUDA). α is the divisor coefficient of the simulation box. The cutoff distance is lower than $1/\beta$ and there are around α^3 non-empty cells in the grid in the Cuda version. The speedup against the CUDA version is shown above the bars for both the build and compute steps, or for the complete run.

4) *GPU Performance comparison*: We then compare the performance of the two GPUs. Results for the A100 are presented in Fig. 9a, 10c, 9e, 9g, and 9i, while results for the RTX8000 are shown in Fig. 9b, 9d, 9f, 9h, and 9j. Both GPUs exhibit comparable performance and a similar initialization-to-computation time ratio. Although the A100 is expected to offer higher raw performance, the OptiX implementation is not fully optimized for this GPU, while the RTX8000 benefits from having more RT cores to accelerate computation.

Additionally, we do not utilize the tensor cores of the A100, which represent a key differentiator between the two GPUs and could help approach theoretical performance.

5) *Comparison between OptiX and CUDA implementations*: Finally, we compare the OptiX triangles (blue), OptiX spheres (green), Custom AABB (orange), Custom sorted AABB (red), and CUDA implementation (green). The OptiX-based implementation is generally slower than the CUDA implementation for β values from 2 to 8, except in a few cases (e.g. $\beta = 2$ and $N = 8$). In these scenarios, the initialization step in the CUDA implementation is negligible, and the computation step dominates. Consequently, avoiding the grid of cells does not yield a significant advantage. For $\beta = 8$ (Fig. 9e and 9f), the CUDA computation step is shorter because the OptiX implementations include not only the interaction computation but also the intersection list computation using rays. Additionally, the memory access pattern in OptiX-based implementations differs, as spatially close particles in memory may not be close in space and may have varying neighbors.

For $\beta = 16$ and $\beta = 32$, the Custom AABB performs better for configurations with few particles per cell (e.g. one or two particles per cell on average in the CUDA version). In these cases, the computation cost is low, and the CUDA version spends most of its time in the initialization step. Furthermore, the CUDA version incurs overhead for empty cells because it constructs a dense grid. For such configurations, the OptiX-based implementations allocate more time to the computation step, making the Custom AABB a suitable choice when the number of particles per cell is low, particularly in scenarios where moving particles require frequent grid rebuilding.

C. Non Uniform Distribution

In this test case, we distribute the particles on the surface of a unit sphere. Consequently, most of the cells on the grid in the CUDA version are empty. With this aim, we use a coefficient α , which can be 8, 16, 32, 64 or 128. For a given α , the simulation grid consists of approximately 3^α non-empty cells. The number of particles N is then calculated as $N = p \times 3^\alpha$, where p represents the average number of particles per cell, taking values of 1, 2, 4, 8, 16, or 32. The cutoff radius is set such that each particle should have approximately $9 \times p$ neighbor particles (it will be in general smaller than $2/\alpha$). In order to facilitate the reproducibility of our test case, we provide in Appendix, the Code 2 that generates the different configurations. All computations are performed in single-precision floating point.

We present the results in Fig. 10. For all configurations, we measured both the initialization step (light color) and the computation step (dark color) for the first configuration. However, for the others, we directly present the total execution

time using a logarithmic scale on the y-axis due to extreme differences.

For the OptiX-based implementation, the initialization step involves building the scene by invoking the OptiX API to create the primitives. For the CUDA version, the initialization step involves constructing the grid of cells. Consequently, in the OptiX-based version, the computation step includes the time spent launching the rays, executing the callback functions, and performing the interactions. In the CUDA version, the computation step corresponds to the kernel time required to compute the interactions.

1) *OptiX spheres*: The OptiX spheres method is consistently the slowest (except for $\alpha = 2$ and a small number of particles). It has a more expensive initialization step compared to the Custom AABB methods and a significantly larger computation step. This is because the built-in IS kernel for spheres is computationally intensive, as it calculates several values, such as the norm of the intersection, that we do not use since we are not rendering an image.

2) *Custom AABB without and with sorting*: For low α , sorting does not provide an advantage; it increases the cost of the initialization step too much relative to the gains in the computation step. However, for high α and a large number of particles, sorting the particles results in a performance improvement. This demonstrates that random access to global memory is so costly that the overhead of sorting the particles is worthwhile. The Custom AABB method with sorting performs comparably to the CUDA version. However, the Custom AABB methods are generally slower than both the CUDA version and the OptiX triangles method.

3) *OptiX triangles*: This strategy is the most efficient. For $\alpha = 8$, its initialization step is comparable in cost to the OptiX spheres and Custom AABB methods. However, its computation step is significantly faster. We attribute this to two primary reasons. First, the built-in IS kernel for triangles is likely much simpler than that for spheres and is probably heavily optimized by NVIDIA, as triangles are the most commonly used primitive in 3D rendering. Second, the AABB bounding boxes around triangles are much smaller than those around spheres or the custom AABB primitives. This enables faster tree traversal and fewer false positives (i.e. cases where a ray intersects a bounding box but not the primitive inside it).

Finally, the CUDA version includes several components in its implementation that rely on algorithms with complexity linear in the number of cells. Since the vast majority of these cells are empty, these steps become highly inefficient. This is particularly evident for $\alpha = 8$, where the initialization step is notably prominent. Moreover, if we were to create a test case with $\alpha = 256$ (not included in the current study), the GPU would run out of memory for the the CUDA version allocating an excessively large grid. This limitation makes alternative approaches based on a tree structure — such as our OptiX-based methods — the only viable options.

VI. CONCLUSION

In this paper, we proposed leveraging ray tracing technology to compute particle interactions within a cutoff distance in a 3D environment. We introduced one method that uses

custom primitives and a custom Intersection Shader (IS), similar to the state-of-the-art, and two methods that use built-in primitives and actual intersection computations. For the latter, we described geometric algorithms to build the interaction list based on ray intersections with spheres or triangles.

Our results indicate that our approach provides a modest advantage in the preprocessing stage by avoiding the construction of a grid of cells. In addition, it is slower than the classical approach during the computation step when dealing with large numbers of uniformly distributed particles. However, when most cells are empty, our approach can provide a significant speedup. Therefore, we believe these methods have the potential to deliver better performance in the future or for specific applications, and we hope our work inspires the community to explore this direction further.

As GPU architectures continue to evolve, with advancements such as NVIDIA's Ada Lovelace architecture featuring third-generation RT cores and AMD's RDNA 3 architecture incorporating second-generation ray-tracing accelerators, we anticipate that algorithms leveraging these capabilities will become increasingly effective and relevant. Consequently, our approach is prepared to benefit from these hardware improvements, enhancing its performance and applicability in future computational scenarios and we hope our work inspires the community to explore this direction further.

We plan to evaluate our approach on other GPUs, such as AMD Radeon, and aim to port our method even in cases where the IS cannot be customized, thanks to our geometric algorithms. In a second step, we will take our best strategy and compare it with state-of-the-art particle interaction solvers that support GPUs.

ACKNOWLEDGMENT

Experiments presented in this paper were carried out using the PlaFRIM experimental test-bed.

REFERENCES

- [1] C. Kutzner, S. Páll, M. Fechner, A. Esztermann, B. L. de Groot, and H. Grubmüller, "More bang for your buck: Improved use of gpu nodes for gromacs 2018," *Journal of Computational Chemistry*, vol. 40, no. 27, pp. 2418–2431, 2019. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/jcc.26011>
- [2] M. Zubair, D. Ranjan, A. Walden, G. Nastac, E. Nielsen, B. Diskin, M. Paterno, S. Jung, and J. H. Davis, "Efficient gpu implementation of automatic differentiation for computational fluid dynamics," in *2023 IEEE 30th International Conference on High Performance Computing, Data, and Analytics (HiPC)*. IEEE, 2023, pp. 377–386.
- [3] B. Rybakina and V. Goryachev, "Parallel algorithms for astrophysics problems," *Lobachevskii Journal of Mathematics*, vol. 39, pp. 562–570, 2018.
- [4] A. Ilievski, V. Zdravski, and M. Gusev, "How cuda powers the machine learning revolution," in *2018 26th Telecommunications Forum (TELFOR)*. IEEE, 2018, pp. 420–425.
- [5] H. Friedrich, J. Günther, A. Dietrich, M. Scherbaum, H.-P. Seidel, and P. Slusallek, "Exploring the use of ray tracing for future games," in *Proceedings of the 2006 ACM SIGGRAPH Symposium on Videogames*, 2006, pp. 41–50.
- [6] NVIDIA, "Nvidia ada gpu architecture," 2023, version 2.1 <https://images.nvidia.com/aem-dam/Solutions/Data-Center/I4/nvidia-ada-gpu-architecture-whitepaper-v2.1.pdf>.
- [7] D. Koschier, J. Bender, B. Solenthaler, and M. Teschner, "Smoothed particle hydrodynamics techniques for the physics based simulation of fluids and solids," *arXiv preprint arXiv:2009.06944*, 2020.

- [8] M. S. Badar, S. Shamsi, J. Ahmed, and M. A. Alam, "Molecular dynamics simulations: concept, methods, and applications," in *Trans-disciplinarity*. Springer, 2022, pp. 131–151.
- [9] D. Algis, B. Bramas, E. Darles, and L. Aveneau, "Efficient GPU Implementation of Particle Interactions with Cutoff Radius and Few Particles per Cell," in *International Symposium on Parallel Computing and Distributed Systems (PCDS2024)*. Singapore, Singapore: IEEE, Sep. 2024. [Online]. Available: <https://inria.hal.science/hal-04621128>
- [10] G. IGN, "Top 5 Games That Make the Best Use of NVIDIA's RTX Technologies," <https://www.ign.com/articles/top-5-games-that-make-the-best-use-of-nvidias-rtx-technologies>, Dec. 2023.
- [11] S. G. Parker, J. Bigler, A. Dietrich, H. Friedrich, J. Hoberock, D. Luebke, D. McAllister, M. McGuire, K. Morley, A. Robison, and M. Stich, "Optix: a general purpose ray tracing engine," *ACM Trans. Graph.*, vol. 29, no. 4, jul 2010. [Online]. Available: <https://doi.org/10.1145/1778765.1778803>
- [12] P. Shirley, I. Wald, T. Akenine-Möller, and E. Haines, "Ray tracing gems: High-quality and real-time rendering with dxr and other apis." Berkeley, CA: Apress, 2019, ch. What is a Ray?, pp. 15–19. [Online]. Available: https://doi.org/10.1007/978-1-4842-4427-2_2
- [13] L. Nylons, M. Harris, and J. Prins, "Fast n-body simulation with cuda," *GPU gems*, vol. 3, pp. 62–66, 2007.
- [14] V. Garcia, E. Debreuve, and M. Barlaud, "Fast k nearest neighbor search using gpu," in *2008 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*. IEEE, 2008, pp. 1–6.
- [15] N. Moenne-Loccoz, A. Mirzaei, O. Perel, R. de Lutio, J. Martinez Esturo, G. State, S. Fidler, N. Sharp, and Z. Gojicic, "3d gaussian ray tracing: Fast tracing of particle scenes," vol. 43, no. 6, Nov. 2024. [Online]. Available: <https://doi.org/10.1145/3687934>
- [16] L. Lindau, "Hardware accelerated ray tracing of particle systems," 2020.
- [17] S. Blyth, "Opticks: Gpu optical photon simulation for particle physics using nvidia® optix™," in *EPJ Web of Conferences*, vol. 214. EDP Sciences, 2019, p. 02027.
- [18] —, "Integration of junos simulation framework with opticks: Gpu accelerated optical propagation via nvidia® optix™," in *EPJ Web of Conferences*, vol. 251. EDP Sciences, 2021, p. 03009.
- [19] A. Keksel, S. Schmidt, D. Beck, and J. Seewig, "Scientific modeling of optical 3d measuring devices based on gpu-accelerated ray tracing using the nvidia optix engine," in *Modeling Aspects in Optical Metrology IX*, vol. 12619. SPIE, 2023, pp. 117–125.
- [20] I. Evangelou, G. Papaioannou, K. Vardis, and A. A. Vasilakis, "Fast radius search exploiting ray tracing frameworks," *Journal of Computer Graphics Techniques (JCGT)*, vol. 10, no. 1, pp. 25–48, February 2021. [Online]. Available: <http://jcgt.org/published/0010/01/02/>
- [21] Y. Zhu, "Rtnn: accelerating neighbor search using hardware ray tracing," in *Proceedings of the 27th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, ser. PPOPP '22. New York, NY, USA: Association for Computing Machinery, 2022, p. 76–89. [Online]. Available: <https://doi.org/10.1145/3503221.3508409>
- [22] S. Zhao, Z. Lai, and J. Zhao, "Leveraging ray tracing cores for particle-based simulations on gpus," *International Journal for Numerical Methods in Engineering*, vol. 124, no. 3, pp. 696–713, 2023. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/nme.7139>

APPENDICES

A. DISCUSSION ON THE SPHERE MODEL

The filtering algorithm presented in Section IV-B is based on the assumption that there is always at least one ray that is intersected once by the sphere and that it is the closest one to the center of the sphere. The following section is dedicated to demonstrating this hypothesis.

In Fig. 11, we show the different possibilities depending on the radius r and ray's length l . As it can be seen, even if the sources located at a distance of C from the target could have their spheres that intersect with the rays when $r > l$, we must set $l = r$ to ensure that the rays will intersect with the

sphere in all cases, especially when the source and target are close.

For the clarity of the proof, a demonstration is first provided in the plane, followed by a generalization to \mathbb{R}^3 .

1) 2D Case: Let us consider a circle of radius r centered at (x_c, y_c) , and a cross⁷ centered at the origin with a length of r . We consider the cases where the circle is located in the first quarter, i.e., $0 \leq x_s \leq r$ and $0 \leq y_s \leq r$, but the demonstration remains valid for other quarters.

The equation of a circle is given by:

$$(x - x_c)^2 + (y - y_c)^2 = r^2. \quad (1)$$

The coordinates of the intersection points of the circle with the axis are given by:

$$\begin{aligned} x_0 &= x_c - \sqrt{r^2 - y_c^2} & \text{and} & & x_1 &= x_c + \sqrt{r^2 - y_c^2} \\ y_0 &= y_c - \sqrt{r^2 - x_c^2} & \text{and} & & y_1 &= y_c + \sqrt{r^2 - x_c^2}. \end{aligned} \quad (2)$$

We provide the Fig. 12 that shows where these points are located on the sphere and the cross.

x_0 and y_0 are the coordinates of the intersection points of the circle that remain on the cross as the circle get away from the origin. On the other hand, x_1 and y_1 are the coordinates of the intersection points of the circle that are the farthest from the origin of the cross and that can potentially be too far to remain on the cross (they will be on the corresponding axis but behind l).

Lemma 1. Given a circle \mathcal{C} and let e_x and e_y be the two segments of the cross of length less than r , then the number of intersections of \mathcal{C} with e_x or with e_y is strictly less than 2.

Proof: As we consider that it is impossible that both x_1 and y_1 exist at the same time, we consider that these two equations cannot be true at the same time

$$x_c + \sqrt{r^2 - y_c^2} \leq r \quad \text{and} \quad y_c + \sqrt{r^2 - x_c^2} \leq r. \quad (3)$$

Which can be simplified as

$$\begin{aligned} x_c &< r - \sqrt{r^2 - y_c^2} & \text{and} & & \sqrt{r^2 - x_c^2} &\leq r - y_c \\ r^2 - x_c^2 &\leq r^2 - 2ry_c + y_c^2 & & & & \\ x_c^2 &> 2ry_c - y_c^2 & & & & \\ x_c &> \sqrt{2ry_c - y_c^2}, & & & & \end{aligned} \quad (4)$$

ending up with

$$\sqrt{2ry_c - y_c^2} < r - \sqrt{r^2 - y_c^2}. \quad (5)$$

For our definition range of $y_c \in [0, r]$, this equation has no solution (see Fig. 13), which confirms that \mathcal{C} will always intersect with the cross at most once for e_x or e_y .

⁷That is, two orthogonal axis-aligned segments of the same size intersecting at their respective centers.

The second statement is to show that the ray that is intersected once is the closest to the center of the sphere.

Lemma 2. Given a circle \mathcal{C} and e one of the two segments of the cross of length less than r , such as the number of intersection of \mathcal{C} with e is equal to 1, then the distance of e with \mathcal{C} is smaller than the distance of the other bar to \mathcal{C} .

Proof: Consider that \mathcal{C} intersects with e_y twice and e_x once, it means that $y_1 < r$ and $x_1 > r$, i.e. $y_c + \sqrt{r^2 - x_c^2} < r$ and $x_c + \sqrt{r^2 - y_c^2} > r$.

We aim to demonstrate that the inequality

$$y_c + \sqrt{r^2 - y_c^2} < x_c + \sqrt{r^2 - x_c^2} \quad (6)$$

holds if and only if $x_c > y_c$.

We start with the inequality:

$$y_c + \sqrt{r^2 - y_c^2} < x_c + \sqrt{r^2 - x_c^2}. \quad (7)$$

Subtracting y_c from both sides, we obtain:

$$\sqrt{r^2 - y_c^2} < x_c - y_c + \sqrt{r^2 - x_c^2}. \quad (8)$$

We can further simplify this to:

$$\sqrt{r^2 - y_c^2} - \sqrt{r^2 - x_c^2} < x_c - y_c. \quad (9)$$

The inequality now compares two quantities: $\sqrt{r^2 - y_c^2} - \sqrt{r^2 - x_c^2}$ and $x_c - y_c$.

- The term $\sqrt{r^2 - y_c^2}$ represents the horizontal distance from the point (x_c, y_c) to the vertical axis.
- The term $\sqrt{r^2 - x_c^2}$ represents the vertical distance from the point (x_c, y_c) to the horizontal axis.

Let's consider the case where $x_c > y_c$:

- If $x_c > y_c$, then $x_c - y_c > 0$.
- Additionally, $\sqrt{r^2 - y_c^2} > \sqrt{r^2 - x_c^2}$ because $y_c < x_c$.

This implies that the term $\sqrt{r^2 - y_c^2} - \sqrt{r^2 - x_c^2}$ is positive, and it is less than $x_c - y_c$, proving that the inequality holds under this condition.

Thus, for the inequality $y_c + \sqrt{r^2 - y_c^2} < x_c + \sqrt{r^2 - x_c^2}$ to hold, it is necessary that $x_c > y_c$. So, the x axis is the closest axis to the center of the sphere and is intersected once.




Fig. 11. 2D Spherical representation of the particles in three different cases with $r > l$, $l > r$ and $r == l$.




Fig. 12. 2D spherical representation illustrating the classification of intersection points between the sphere and the cross.




Fig. 13. Plot of the equation $\sqrt{2ry_c - y_c^2} - r + \sqrt{r^2 - y_c^2}$, for $r = 1$.

2) *3D Case:* To convert this proof from 2D to 3D, we need to extend the concepts from the circle and cross to a sphere and a three dimensional cross.⁸ Consider a sphere with radius r centered at (x_c, y_c, z_c) in 3D space, and a cross (or coordinate axes) centered at the origin with each axis extending from $-l$ to l . We are interested in analyzing the intersection of the sphere with the axes, focusing particularly on the first octant where $0 \leq x_c \leq l$, $0 \leq y_c \leq l$, and $0 \leq z_c \leq l$.

The equation of the sphere is given by:

$$(x - x_c)^2 + (y - y_c)^2 + (z - z_c)^2 = r^2. \quad (10)$$

The coordinates of the intersection points of the sphere with the axes are found by setting two of the coordinates to zero in the sphere's equation:

- Intersection with the x -axis (set $y = 0$ and $z = 0$):

$$x = x_c \pm \sqrt{r^2 - y_c^2 - z_c^2} \quad (11)$$

⁸That is, three axis aligned orthogonal segments of the same size intersecting at their respective centers.

- Intersection with the y -axis (set $x = 0$ and $z = 0$):

$$y = y_c \pm \sqrt{r^2 - x_c^2 - z_c^2} \quad (12)$$

- Intersection with the z -axis (set $x = 0$ and $y = 0$):

$$z = z_c \pm \sqrt{r^2 - x_c^2 - y_c^2} \quad (13)$$

Let's denote the intersection points on the positive half of the axes as:

- $x_1 = x_c + \sqrt{r^2 - y_c^2 - z_c^2}$
- $y_1 = y_c + \sqrt{r^2 - x_c^2 - z_c^2}$
- $z_1 = z_c + \sqrt{r^2 - x_c^2 - y_c^2}$

We need to analyze whether these points lie within the bounds of the cross, i.e. whether $x_1 \leq l$, $y_1 \leq l$, and $z_1 \leq l$.

Assume that one of these coordinates, say x_1 , exceeds l . This would mean that the intersection does not lie on the cross, i.e. $x_c + \sqrt{r^2 - y_c^2 - z_c^2} > l$.

Similarly, for y_1 and z_1 , we require that:

$$y_c + \sqrt{r^2 - x_c^2 - z_c^2} > l \quad \text{or} \quad z_c + \sqrt{r^2 - x_c^2 - y_c^2} > l. \quad (14)$$

These conditions cannot all be true simultaneously for x_c , y_c , and z_c within the defined range, similar to the 2D case. Thus, a sphere will intersect the cross at most once per axis.

Next, we determine the axis closest to the sphere's center. If $x_c > y_c > z_c$, we aim to prove that the intersection on the x -axis occurs first (i.e. is the smallest).

Starting with:

$$x_c + \sqrt{r^2 - y_c^2 - z_c^2} < y_c + \sqrt{r^2 - x_c^2 - z_c^2} \quad (15)$$

and $x_c + \sqrt{r^2 - y_c^2 - z_c^2} < z_c + \sqrt{r^2 - x_c^2 - y_c^2}$.

These can be simplified, following similar steps as in the 2D case:

$$\sqrt{r^2 - y_c^2 - z_c^2} - \sqrt{r^2 - x_c^2 - z_c^2} < y_c - x_c \quad (16)$$

and $\sqrt{r^2 - y_c^2 - z_c^2} - \sqrt{r^2 - x_c^2 - y_c^2} < z_c - x_c.$

The argument follows that since $x_c > y_c > z_c$, the inequalities hold true, confirming that the x -axis is the closest, and thus it is intersected first.

The 3D extension of the proof shows that a sphere intersects each axis of a coordinate cross at most once, and the axis closest to the sphere's center (in the order of $x_c > y_c > z_c$) will have the intersection point closest to the origin.

B. CONVERSION FROM PARTICLES' POSITIONS TO TRIANGLES

```

1  for(int i = 0; i < nbPoints; i++)
2  {
3      const float3 point = points[i];
4      std::array<float3, 8> corners;
5      for(int idxCorner = 0; idxCorner < 8; ++idxCorner){
6          corners[idxCorner].z = point.z + (idxCorner&1 ? ...
7              (cutoffRadius/2)+epsilon : (-cutoffRadius/2)-epsilon);
8          corners[idxCorner].y = point.y + (idxCorner&2 ? ...
9              (cutoffRadius/2)+epsilon : (-cutoffRadius/2)-epsilon);
10         corners[idxCorner].x = point.x + (idxCorner&4 ? ...
11             cutoffRadius/2 : -cutoffRadius/2);
12     }
13     vertices.push_back(corners[0]);
14     vertices.push_back(corners[1]);
15     vertices.push_back(corners[3]);
16
17     vertices.push_back(corners[0]);
18     vertices.push_back(corners[2]);
19     vertices.push_back(corners[3]);
20
21     vertices.push_back(corners[4]);
22     vertices.push_back(corners[5]);
23     vertices.push_back(corners[7]);

```

```

22     vertices.push_back(corners[4]);
23     vertices.push_back(corners[6]);
24     vertices.push_back(corners[7]);
25 }
26

```

Code 1: Triangles generation from particles's positions.

C. GENERATION OF PARTICLES ON A SPHERE

```

1  const int MaxParticlesPerCell = 32;
2  const int MaxBoxDiv = 128;
3  for(int boxDiv = 2; boxDiv <= MaxBoxDiv; boxDiv *= 2){
4      const int nbBoxes = boxDiv*boxDiv*boxDiv;
5      for(int nbParticles = nbBoxes; nbParticles <= ...
6          nbBoxes*MaxParticlesPerCell; nbParticles *= 2){
7          const double particlePerCell = ...
8          double(nbParticles)/double(nbBoxes);
9          const double expectedNbNeighbors = 9*particlePerCell;
10         const double coef = 1. - ...
11         ((2*expectedNbNeighbors)/nbParticles);
12         const double validCoef = std::min(1.0, std::max(-1.0, coef));
13         const double sphereRadius = acos(validCoef);
14
15         // The following is only used if we need to build a grid ...
16         of cell
17         const float boxWidth = std::ceil(2.0 / sphereRadius) * ...
18         sphereRadius;
19         const int gridDim = boxWidth/sphereRadius;
20         const float cellWidth = boxWidth/gridDim;
21         .....
22         auto generateRandomParticle() {
23             double theta = 2.0 * M_PI * ((double)rand() / RAND_MAX); // ...
24             Random angle between 0 and 2PI
25             double phi = acos(1.0 - 2.0 * ((double)rand() / RAND_MAX)); // ...
26             // Random angle between 0 and PI
27
28             // Convert spherical coordinates to Cartesian coordinates
29             double x = sin(phi) * cos(theta);
30             double y = sin(phi) * sin(theta);
31             double z = cos(phi);
32
33             return Particle(x, y, z);
34         }

```

Code 2: Non-uniform test case generation.