

Ray Tracing: Acceleration Techniques

1 The Complexity Problem

Ray tracing is computationally expensive. To understand why, consider a typical render:

- **Resolution:** 1024×1024 pixels ($\approx 10^6$ pixels); cast 10^6 rays
- **Scene:** 1 million (10^6) triangles; do 10^{12} ray-triangle intersections
- **Samples:** 10 rays per pixel for antialiasing; 10^{13} intersections
- **Lighting:** 10 lights (shadow rays!); 10^{14} intersections

Now, for the processing power.


- **Cores:** 10 cores; 10^{13} intersections/core
- **Cost:** 100 clock cycles/intersection; 10^{15} clock cycles
- **Speed:** 3GHz ; 3.3×10^5 seconds
- 10^2 hours
- 4 days

That is not practical.

2 Bounding Volumes


The most fundamental acceleration technique is the use of Bounding Boxes. Instead of testing a ray against every triangle in a complex object, we first test it against a simple proxy volume.

- If the ray misses the bounding box, it cannot possibly hit the object inside.
- If the ray hits the bounding box, we then perform the more expensive tests on the geometry within.




3 Spatial Hierarchies

Organize bounding volumes into hierarchies. Then, check recursively. This allows us to prune out large portions of the scene for any given ray.



Just sorting from left to right does not actually work well. Instead, at each level of the hierarchy, we split along a different direction. This tends to keep the boxes small in all directions.



4 Grid-based data structure

An alternative data structure is to use a grid, storing a list of pointers to objects in each grid cell. To ray trace the scene, walk the grid cells as they are encountered along the ray. Pointers to the same object need not be checked multiple times if they occur in many cells. Traversal stops as soon as the an intersection is identified in the current cell. In the figure below, the objects are **B**, **G**, **R**, and **O**. Traversal ends at cell 6, where the first intersection occurs.

