





# Memory-Efficient GPU Volume Path Tracing of AMR Data Using the Dual Mesh—Supplementary Document


Stefan Zellmann<sup>1</sup> , Qi Wu<sup>2</sup> , Kwan-Liu Ma<sup>2</sup> , and Ingo Wald<sup>3</sup>   
<sup>1</sup>University of Cologne <sup>2</sup>University of California - Davis <sup>3</sup>NVIDIA

## 1. Additional Color Images

This document provides supplementary material to our paper [ZWMW23]. We present several high-resolution color images that would not fit into the paper due to space limitations. We use presentations similar to Fig. 1 from the paper to convey the spatial structure of the data. Fig. 1 shows the NASA Exajet data set from the front; Fig. 2 shows Exajet from a rear view. The structures near the wing and fuselage are finely tessellated with tiny dual cells. Fig. 3 shows the LANL meteor impact data set; note the finely tessellated ocean surface. We use a clip plane to reveal and further convey the structure in Fig. 4. Fig. 5 shows the NASA Landing Gear, again with a visualization inspired by Fig. 1 from the paper. The data set suffers from the “teapot in a stadium” problem, as can be seen in Fig. 6.

## 2. Additional Evaluation

In the paper we mention that we use gridlets of size  $8^3$  cells that store  $9^3$  voxels. Fig. 7 presents an evaluation of gridlet size vs. memory consumption, both manually computed and measured with `nvidia-smi`. Smaller gridlets have more ghost cells as scalars at gridlet boundaries are replicated; bigger gridlets contain more empty cells. Bigger gridlets also result in a more shallow BVH; we still observed negative returns in both memory as well as rendering performance when scaling beyond gridlet sizes of  $8^3$ .



**Figure 7:** Gridlet size vs. memory and rendering performance. We use a gridlet size of  $8^3$  cells ( $9^3$  scalars); gridlet sizes beyond that show negative returns both in memory consumption (“data”, manually computed and “final”, measured with `nvidia-smi`) and in rendering performance (in frames/second, FPS).

Note that the differences in frames/second are still subtle: always

approximately in the range of 7-8 FPS for the TAC Cloud, and 4.5-5 FPS for LANL Impact.


We use a bilinear face test for boundary stitching elements. For that we ported the implementation from OpenVKL [Int] for pyramids, wedges, and hexes. We use ten Newton iterations to refine the result. In Fig. 8 we present a visual comparison, including difference images and contrast-enhanced images to highlight the subtle differences.

## 3. Code Listings


We provide code snippets in C++ pseudo code that represent gridlet construction (Algorithm 1) Figs. 9 and 10. This implementation assumes that a hashmap is used to represent the macrocells of the virtual grid, i.e., macrocells that are (and stay) empty are never created in memory. Hashmaps are not necessarily the most sensible choice to use when focusing on performance; however, they allow for a very simple implementation of Algorithm 1. In Fig. 11 we provide CUDA pseudo code to sample the gridlet element type in the shader or ray tracing program.

## References


- [ANA\*20] ANDERSSON P., NILSSON J., AKENINE-MÖLLER T., OSKARSSON M., ÅSTRÖM K., FAIRCHILD M. D.: FLIP: A Difference Evaluator for Alternating Images. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 3, 2 (2020), 15:1–15:23. 5
- [Int] INTEL CORPORATION: High Performance Volume Kernels. Available at <https://www.openvkl.org/>, Accessed: 11 November 2022. 1
- [ZWMW23] ZELLMANN S., WU Q., MA K.-L., WALD I.: Memory-Efficient GPU Volume Path Tracing of AMR Data Using the Dual Mesh. *Computer Graphics Forum (Proc. EuroVis)* (2023). 1




**Figure 1:** Exajet front view. Top: shaded (top-right) and voxels+stitching elements (bottom-left). Bottom: gridlets colored by primitive ID.




**Figure 2:** Exajet back view. Top: shaded (top-right) and voxels+stitching elements (bottom-left). Bottom: gridlets colored by primitive ID.




**Figure 3:** LANL impact. Left: shaded (top-right) and voxels+stitching elements (bottom-left). Right: gridlets colored by primitive ID.




**Figure 4:** LANL impact. Left: shaded (top-right) and voxels+stitching elements (bottom-left). Right: gridlets colored by primitive ID. We use a clip plane to reveal the inner structure of the ocean and plume.



**Figure 5:** NASA Landing Gear. Left: shaded (top-right) and voxels+stitching elements (bottom-left). Right: gridlets colored by primitive ID.



**Figure 6:** NASA Landing Gear, zoomed out. Left: showing gridlets colored by primitive ID. Right: same, but with a different transfer function. The gear itself is seen at a distance. This data set suffers from the “teapot in a stadium” problem.



**Figure 8:** We use sampling tests that model twisted stitching elements (a) with bilinear faces (b). Element tests that approximate that with planar faces exhibit subtle artifacts (c). In (d) we show a difference image generated with FLIP [ANA\*20]. (e) shows a contrast-enhanced region of interest of (b). (f) shows a contrast-enhanced region of interest of (c).

```

1 // static macrocell size
2 const static int3 g_mcSize{16,16,16};
3
4 // compute macrocell ID from voxel
5 int3 mcID(Voxel voxel) {
6     return voxel.lower/g_mcSize;
7 }

```

**Figure 9:** Mapping voxel coordinates to macrocell IDs.

```

1 void makeGrids(int level) {
2     // map from voxel corners (int3)
3     // to macrocells (3D box, integer coordinates)
4     map<int3,box3i> macrocells;
5
6     // macrocells are invalid initially; lazily
7     // "activate" only when cubes overlap
8     for (auto cube : cubes[level]) {
9         Voxel vox(cube,level);
10        macrocells[mcID(vox)].extend(cube.bounds());
11    }
12
13    // Iterate over all *active* macrocells;
14    // these become "gridlets"
15    for (auto mc : macrocells) {
16        int3 numVoxels = mc.second.size();
17        // gridlets store data at the cell corners:
18        int numScalars = (numVoxels.x+1)
19                        * (numVoxels.y+1)
20                        * (numVoxels.z+1);
21
22        Gridlet gridlet = {
23            .lower = mc.second.lower,
24            .level = level,
25            .numScalars = numScalars,
26            .scalarIDs = int[numScalars]
27        };
28
29        // Add to gridlet list
30        g_gridlets[level].add(gridlet);
31    }
32 }

```

**Figure 10:** Algorithm to cluster voxels to gridlets using a virtual uniform grid of macrocells.

```

1 bool sample(Gridlet g, float3 pos, float *value){
2     box3f bounds = g.bounds();
3     if (bounds.contains(pos)) {
4         // gridlets store data at the corners
5         int3 numScalars = g.dims+1;
6
7         // compute box around position in raster
8         // coords on the gridlet's level
9         int3 imin = (pos-bounds.lower)/(1<<g.level);
10        int3 imax = min(imin+1,numScalars-1);
11
12        // data values at the box corners
13        float values[8] = {
14            g.scalars[int3(imin.x,imin.y,imin.z)],
15            g.scalars[int3(imax.x,imin.y,imin.z)],
16            g.scalars[int3(imin.x,imax.y,imin.z)],
17            g.scalars[int3(imax.x,imax.y,imin.z)],
18            g.scalars[int3(imin.x,imin.y,imax.z)],
19            g.scalars[int3(imax.x,imin.y,imax.z)],
20            g.scalars[int3(imin.x,imax.y,imax.z)],
21            g.scalars[int3(imax.x,imax.y,imax.z)]
22        };
23
24        if (all(values) != NAN) {
25            // trilinear interpolation
26            float3 p = (pos-bounds.lower)/(1<<g.level);
27            float3 frac = p-imin;
28            *value = trilerp(values,frac);
29
30            return true;
31        }
32    }
33
34    // either missed bounds, or hit an empty cell
35    return false;
36 }

```

**Figure 11:** Gridlet point sampling function.