Supplemental Material for
"Profiling and Visualizing GPU Memory Access and Cache Behavior of Ray Tracers"

Max von Buelow[1] ⓘ, Kai Riemann[2] ⓘ, Stefan Guthe[1,3] ⓘ and Dieter W. Fellner[1,3,4] ⓘ

[1]Technical University of Darmstadt, Germany
[2]Independent Researcher, Germany
[3]Fraunhofer IGD, Germany
[4]Graz University of Technology, Institute of Computer Graphics and Knowledge Visualization, Austria

This supplemental material is structured as follows:

In section A, we shortly introduce the meshes used in this supplemental material. Figure 1 demonstrates all eight frames from the time series visualization (fig. 1 of the main paper) and its corresponding single-frame visualization. Additionally, we provide the same time series visualization for two further meshes: San Miguel and the Stanford Bunny. The effects of multiple ray tracer configurations on the Dragon mesh are shown in fig. 2. Compared to fig. 6 of the main paper, we provide a differencing view of the geometry hit rates between our default configuration and all remaining ones. Additionally, we provide a visualization of the bare triangle access rates as well as a merged BVH and geometry visualization. Figure 3 demonstrates these effects using the additional San Miguel mesh. Section B shortly describes the differences between the Dragon and San Miguel configuration comparison.

A. Mesh Descriptions

The XYZ RGB Asian Dragon and the Crytek Sponza meshes were already introduced in section 6.1 of the main paper. Unlike the main paper, this supplemental material includes two further meshes. San Miguel by GUILLERMO M. LEAL LLAGUNO is a modeled scene showing the atrium of a house with varying triangle resolutions. Similar to the Sponza mesh, we placed the camera inside the house pointing towards the atrium. The San Miguel mesh contains 10 million triangles. The Stanford Bunny is a small 3D-scanning and mesh-zippering data set from the Stanford scanning repository totaling 69 451 triangles.

B. Configuration Comparison

We briefly describe the differences between the San Miguel mesh and the Dragon mesh in this section. The first obvious difference is that the median split (fig. 3b) leads to a less efficient hit rate compared to the Dragon mesh on both the geometry and the BVH. This can be explained by the fact that the San Miguel mesh is less compact, making a static cut on the median less efficient. A second observation is that BVH hit rates are more descriptive. The median-split visualization and the visualization of the *if-if* approach show more nodes with low hit rates. Compared to the Dragon mesh, we do not recognize any patterns on the mesh, which can be explained by extremely low triangle resolutions on flat areas simply blurring out that effect and the great number of submeshes. The visualization of the bare triangle access rate is also more varying on the San Miguel mesh, which is again caused by its multi-resolution triangle representation.

(a) *Dragon (global scan)*

(b) *Dragon (SM-wise scan)*

(c) *San Miguel*

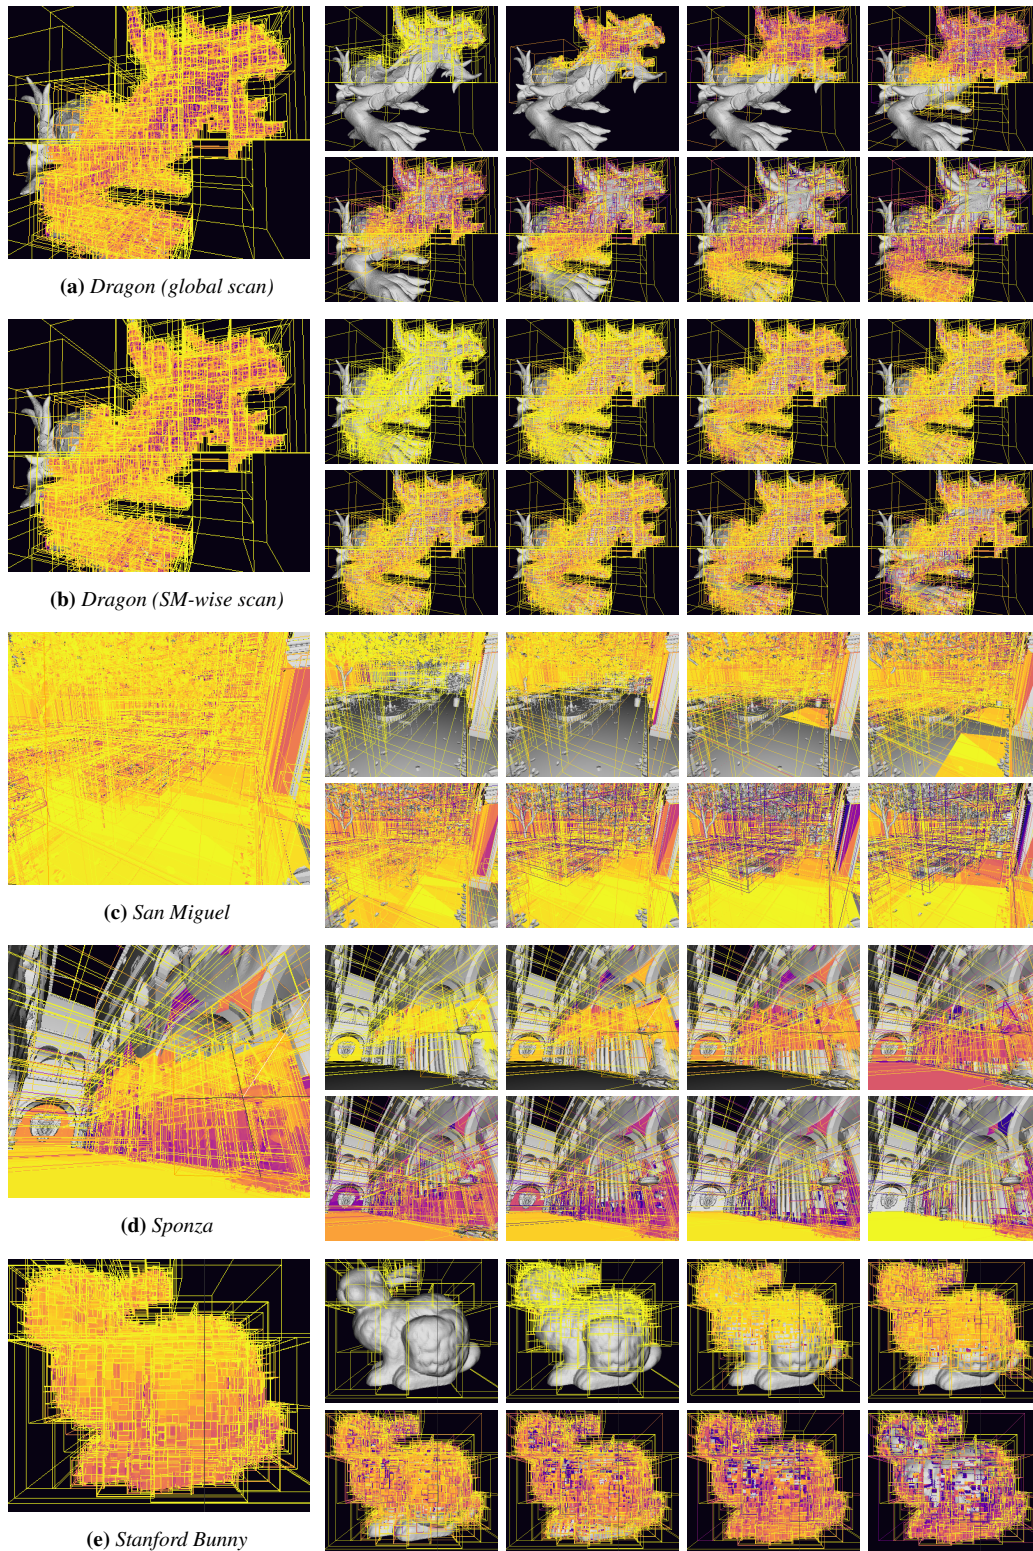**(d)** *Sponza*

(e) *Stanford Bunny*

Figure 1: *Time series visualization on the Dragon with global scanline scheduling (a) and SM-wise scheduling (b), on San Miguel (c), on Sponza (d), and on the Stanford Bunny (e). The large view on the left-hand side shows the single-frame visualization, whereas the small views on the right-hand side display the progression over eight time frames.*
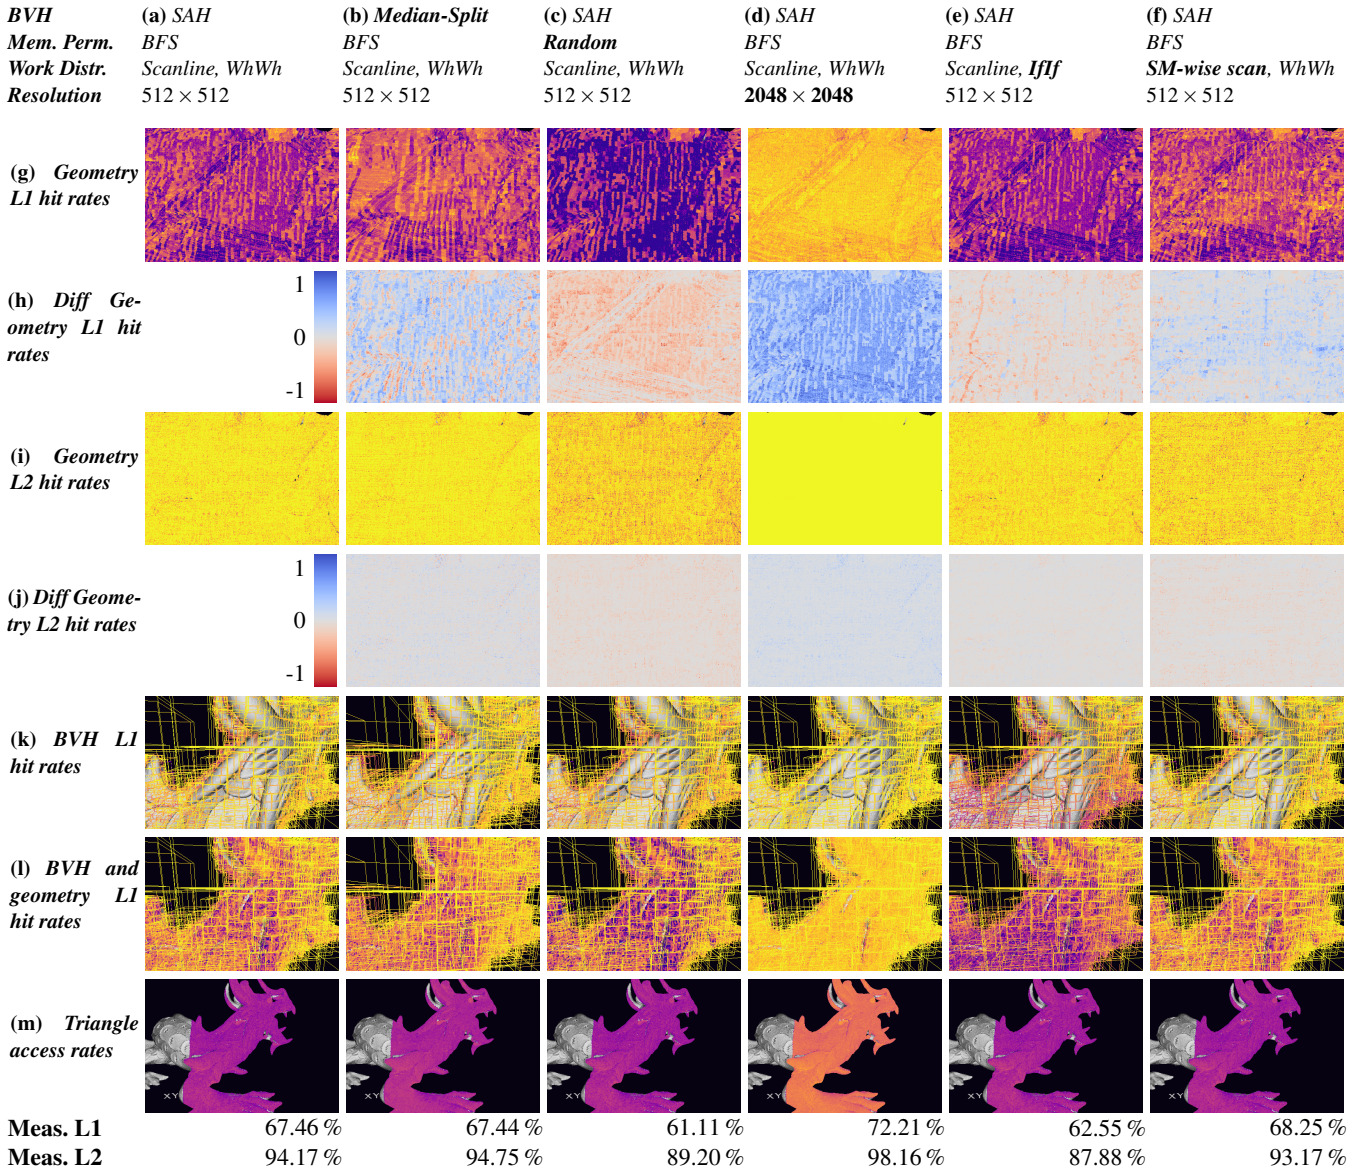
BVH	**(a)** *SAH*	**(b)** *Median-Split*	**(c)** *SAH*	**(d)** *SAH*	**(e)** *SAH*	**(f)** *SAH*
Mem. Perm.	*BFS*	*BFS*	*Random*	*BFS*	*BFS*	*BFS*
Work Distr.	*Scanline, WhWh*	*Scanline, WhWh*	*Scanline, WhWh*	*Scanline, WhWh*	*Scanline, IfIf*	*SM-wise scan, WhWh*
Resolution	512 × 512	512 × 512	512 × 512	**2048 × 2048**	512 × 512	512 × 512

(g) *Geometry L1 hit rates*						
(h) *Diff Geometry L1 hit rates*						
(i) *Geometry L2 hit rates*						
(j) *Diff Geometry L2 hit rates*						
(k) *BVH L1 hit rates*						
(l) *BVH and geometry L1 hit rates*						
(m) *Triangle access rates*						
Meas. L1	67.46 %	67.44 %	61.11 %	72.21 %	62.55 %	68.25 %
Meas. L2	94.17 %	94.75 %	89.20 %	98.16 %	87.88 %	93.17 %

Figure 2: *Visualizations using the Asian Dragon. This row presents the same zoomed-in region of the geometry L1 hit rate as in the main paper (g) followed by an difference view (h) that represents the difference between our standard configuration (a) and all other configurations. Blue colors highlight areas where hit rates are higher, red colors represent areas where hit rates are lower. The color map is included at the left-hand side of the row. The subsequent two rows (i) (j) present the same for the L2 hit rates. The following line (k) shows the same region of the BVH L1 hit rate visualization as in the main paper and the next row (l) merged with geometry L1 hit rates. The final row (m) shows bare triangle access rates.*
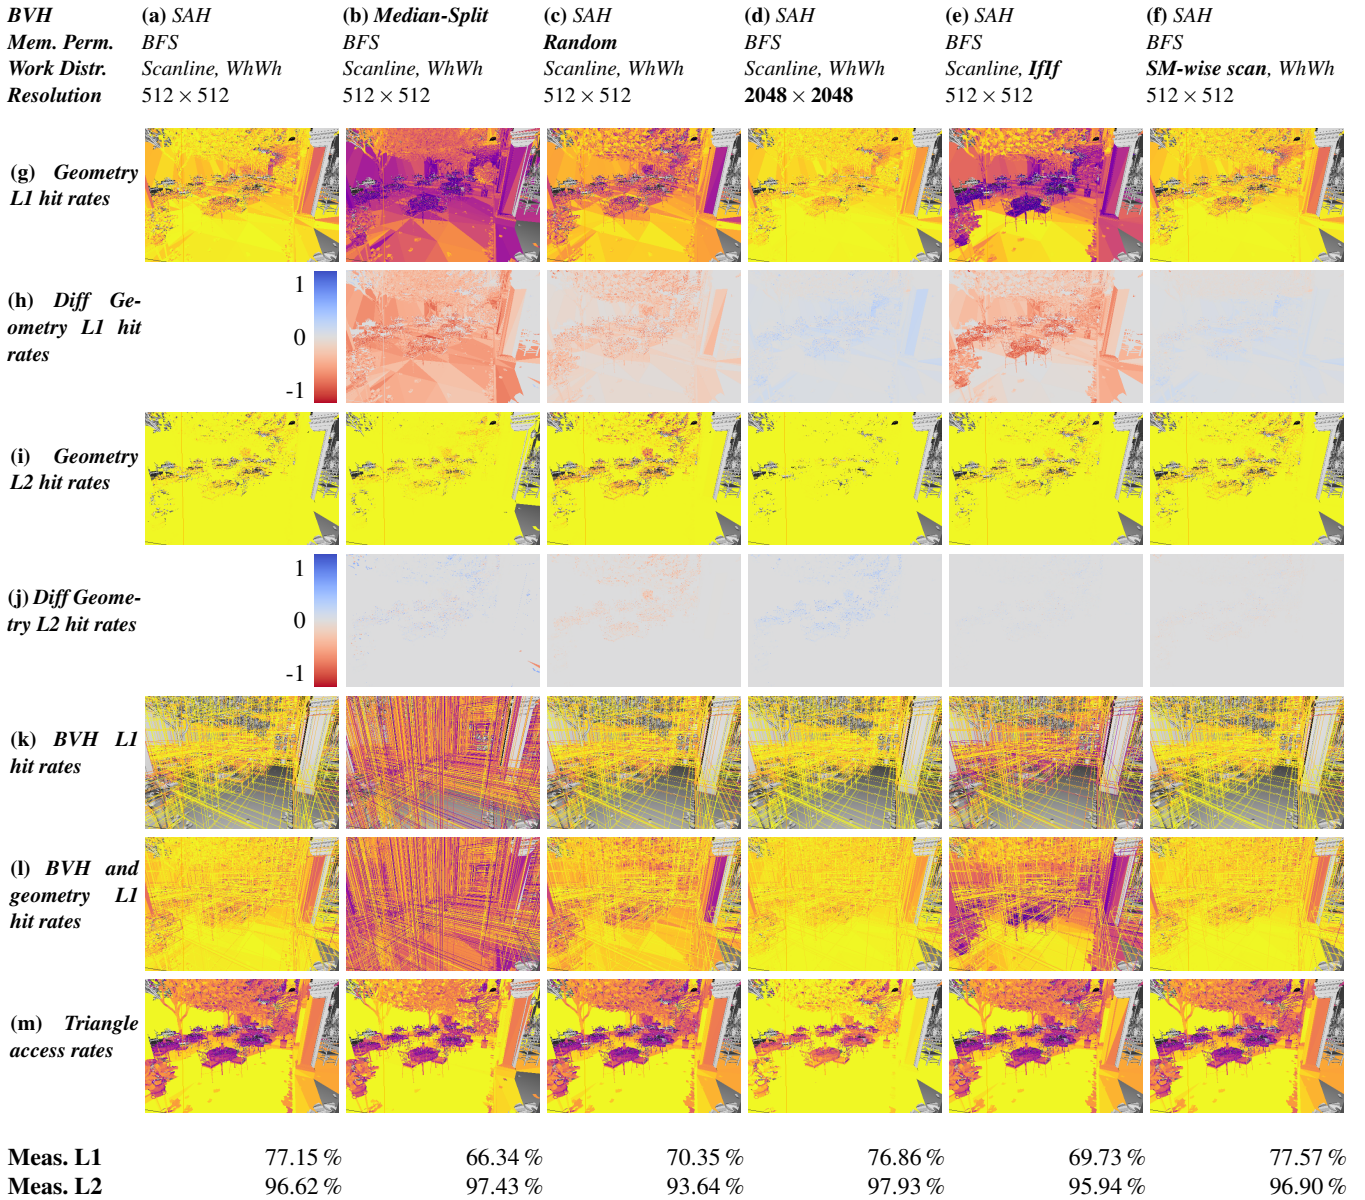
BVH	(a) *SAH*	(b) ***Median-Split***	(c) *SAH*	(d) *SAH*	(e) *SAH*	(f) *SAH*
Mem. Perm.	*BFS*	*BFS*	***Random***	*BFS*	*BFS*	*BFS*
Work Distr.	*Scanline, WhWh*	*Scanline, WhWh*	*Scanline, WhWh*	*Scanline, WhWh*	*Scanline, **IfIf***	***SM-wise scan**, WhWh*
Resolution	512 × 512	512 × 512	512 × 512	**2048 × 2048**	512 × 512	512 × 512

	(a)	(b)	(c)	(d)	(e)	(f)
Meas. L1	77.15 %	66.34 %	70.35 %	76.86 %	69.73 %	77.57 %
Meas. L2	96.62 %	97.43 %	93.64 %	97.93 %	95.94 %	96.90 %

Figure 3: *Visualizations using the San Miguel mesh. See fig. 2 for a description of the matrix.*