Hybrid Voxel Formats for Efficient Ray Tracing

RUSSEL ARBORE, JEFFREY LIU, AIDAN WEFEL, STEVEN GAO, ERIC SHAFFER
UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

Problem

Voxels are a geometric representation useful for rendering volumes, multi-resolution models, and indirect lighting effects. The memory consumption of uncompressed voxel volumes scales cubically with their resolution, making it impractical to store and render them at high resolutions.

Our work explores the design of alternative voxel storage formats to optimize tradeoffs between storage size and rendering performance.

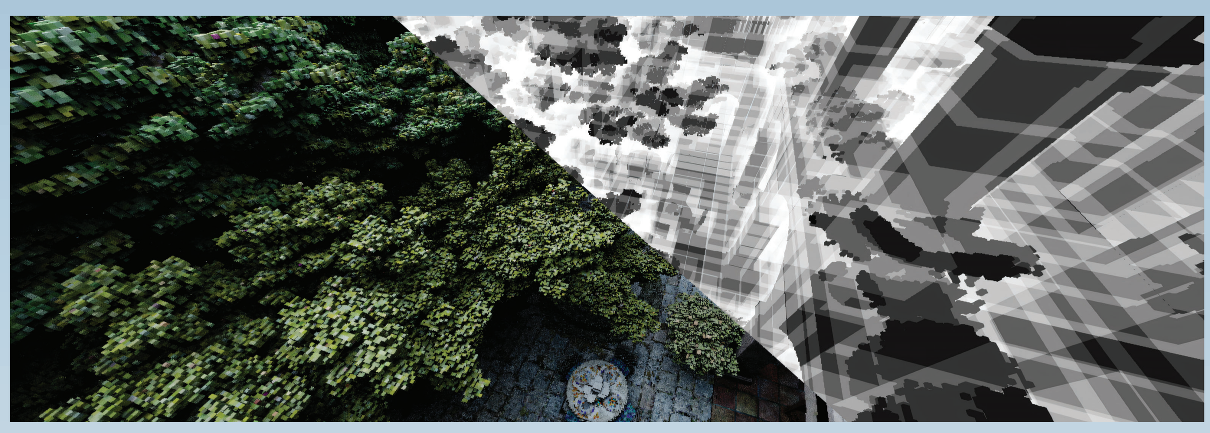

Voxel Formats and Previous Work

Our method builds off of previous work in optimizing voxel storage and ray tracing.

- Raw Grids - store voxels directly as uniform grids in memory [1, 2]
- Sparse Voxel Octrees (SVO) - partition space recursively to compress homogeneous areas [3]
- Sparse Voxel Directed Acyclic Graphs (SVDAG) - modify an SVO by deduplicating nodes [4]
- Distance Fields (DF) - accelerate raytracing by storing a distance to nearest nonempty voxel

Each format is laid out differently in memory, resulting in tradeoffs between storage and ray intersection performance.

Hybrid Voxel Formats

We present a formulation of "hybrid" formats, wherein each level of a hierarchical format can feature a different structure. We show that hybrid voxel formats can achieve Pareto optimal trade-offs between memory consumption and rendering performance.

References

[1] John Amanatides and Andrew Woo. A Fast Voxel Traversal Algorithm for Ray Tracing. (EuroGraphics 1987)
[2] Cyril Crassin, Fabrice Neyret, Sylvain Lefebvre, and Elmar Eisemann. GigaVoxels: Ray-Guided Streaming for Efficient and Detailed Voxel Rendering. (I3D 2009)
[3] Samuli Laine and Tero Karras. Efficient sparse voxel octrees. (I3D 2010)
[4] Viktor Kämpe, Erik Sintorn, and Ulf Assarsson. High resolution sparse voxel DAGs. (SIGGRAPH 2013)
[5] Morgan McGuire. Computer Graphics Archive. (https://casual-effects.com/data)

Author Contact

Russel Arbore - rarbore2@illinois.edu
Jeffrey Liu - jliu179@illinois.edu
Aidan Wefel - awefel2@illinois.edu
Steven Gao - hongyig3@illinois.edu
Eric Shaffer - shaffer1@illinois.edu

Method

Hybrid formats can be built by composing base formats hierarchically. A single voxel is an "upper" level format is an entire sub-volume, where the geometry of the sub-volume is the data stored in the "lower" level format.

Format	Parameters	Description
Raw	(W, H, D)	Raw grid of (W × H × D) voxels
DF	(W, H, D, M)	Grid of (W × H × D) voxels with maximum Hamming distance M to nearest non-empty voxel
SVO	(L)	Sparse Voxel Octree with maximum depth L
SVDAG	(L)	Sparse Voxel Directed Acyclic Graph with maximum depth L

Format descriptions and associated parameters

2D slice of an **SVO(3) Raw(4, 4, 4) Raw(8, 8, 8)**

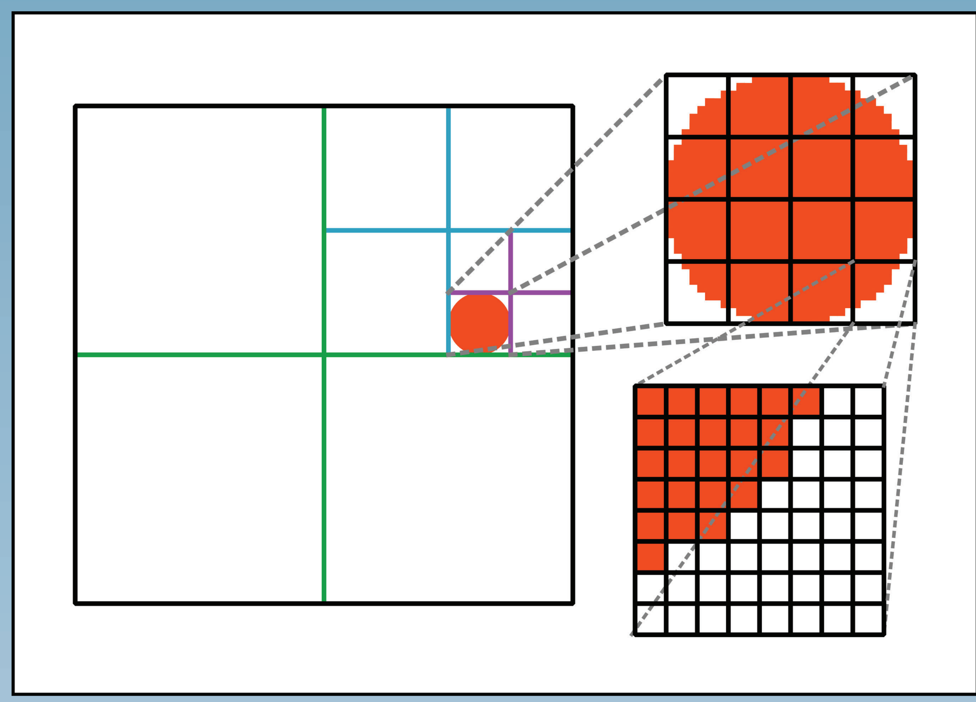

We use a metaprogramming system to specify hybrid formats with a sequence of format descriptions, which automatically generates C++ construction and GLSL intersection code. The format descriptor above is an example input.

Results

We voxelize 4 models from the Computer Graphics Archive [5] and convert them into a set of 2048^3 hybrid voxel formats. We then measure the average frametime of launching 2,073,600 primary rays.

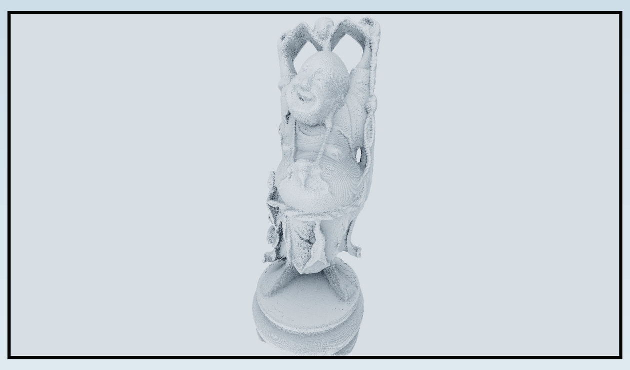

Buddha © 1996 Stanford University

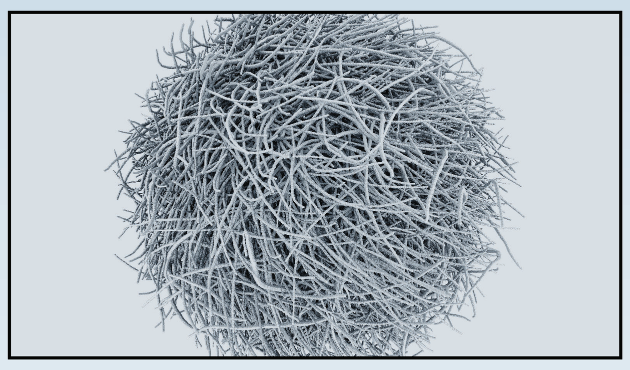

Hairball © 2010 NVIDIA Research

San Miguel © Guillermo M. Leal Llaguno

Sponza © 2002 Marko Dabrovic

Label	Format
I	DF(16, 16, 16, 6) DF(8, 8, 8, 6) SVDAG(4)
II	DF(16, 16, 16, 6) Raw(8, 8, 8) SVDAG(4)
III	Raw(16, 16, 16) Raw(8, 8, 8) SVDAG(4)
IV	Raw(16, 16, 16) SVO(3) SVDAG(4)
V	Raw(16, 16, 16) Raw(16, 16, 16) Raw(8, 8, 8)
VI	DF(16, 16, 16, 6) SVO(7)
VII	DF(16, 16, 16, 6) SVDAG(7)
VIII	DF(64, 64, 64, 6) SVO(5)
IX	DF(64, 64, 64, 6) SVDAG(5)
X	Raw(16, 16, 16) SVO(7)
XI	Raw(16, 16, 16) SVDAG(7)
XII	Raw(64, 64, 64) SVO(5)
XIII	Raw(64, 64, 64) SVDAG(5)
XIV	Raw(8, 8, 8) SVDAG(8)
XV	Raw(256, 256, 256) SVDAG(3)
XVI	SVO(7) SVDAG(4)
XVII	SVO(5) SVDAG(6)
XVIII	SVO(3) SVDAG(8)
XIX	SVO(11)
XX	SVDAG(11)

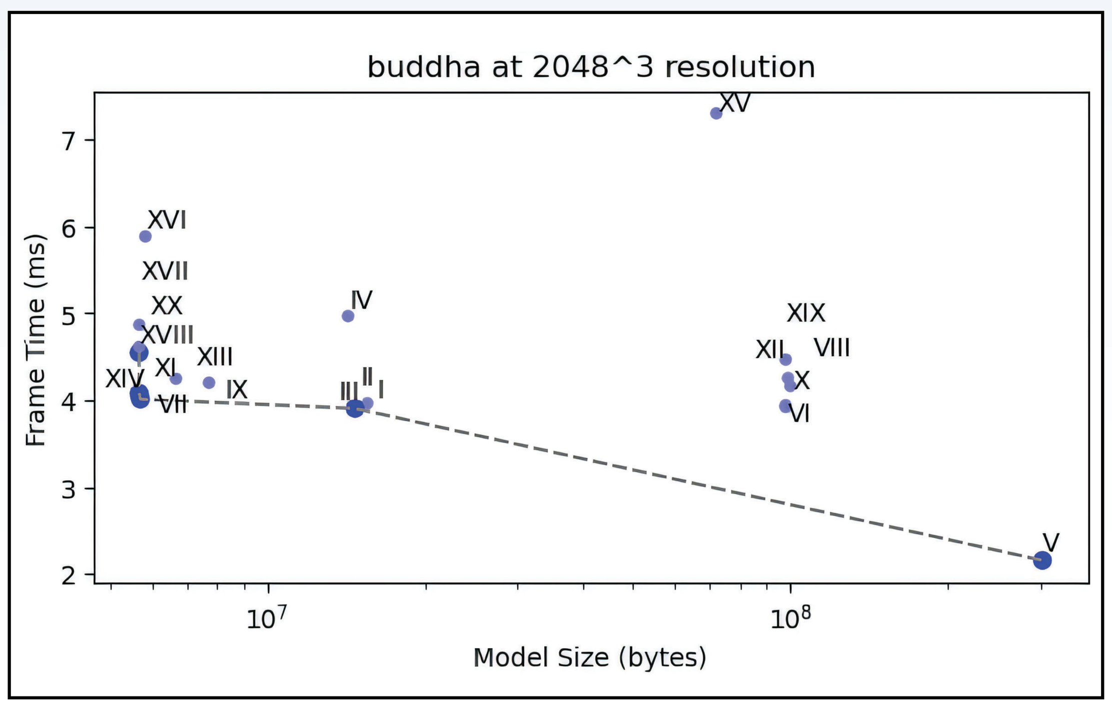

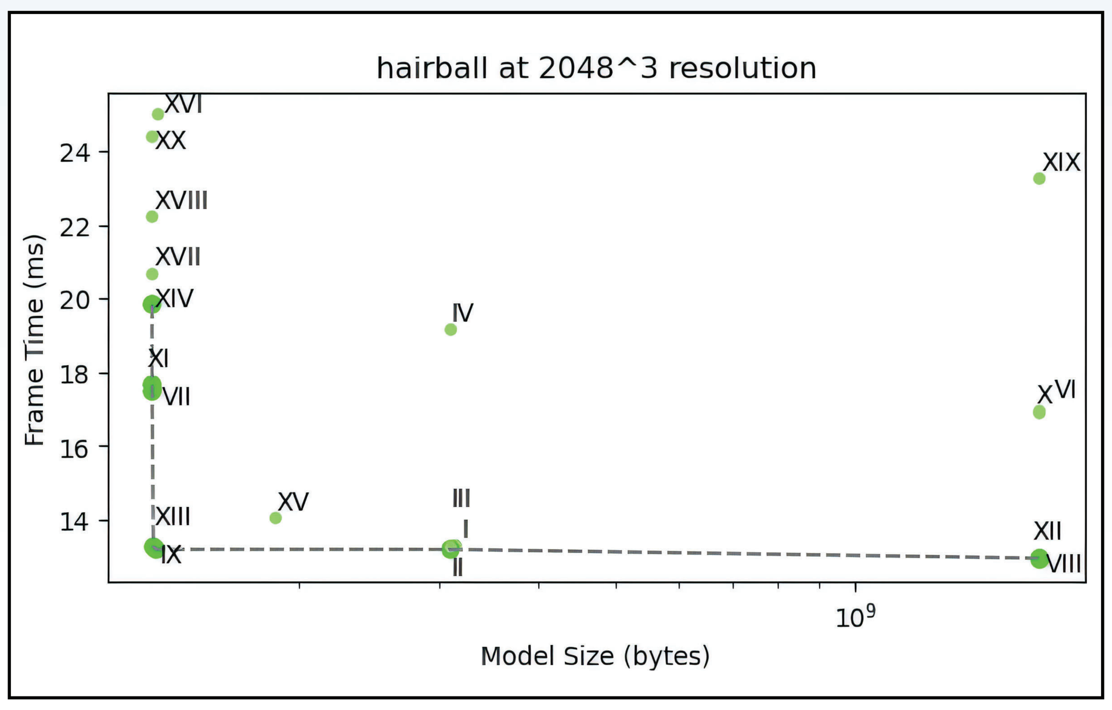

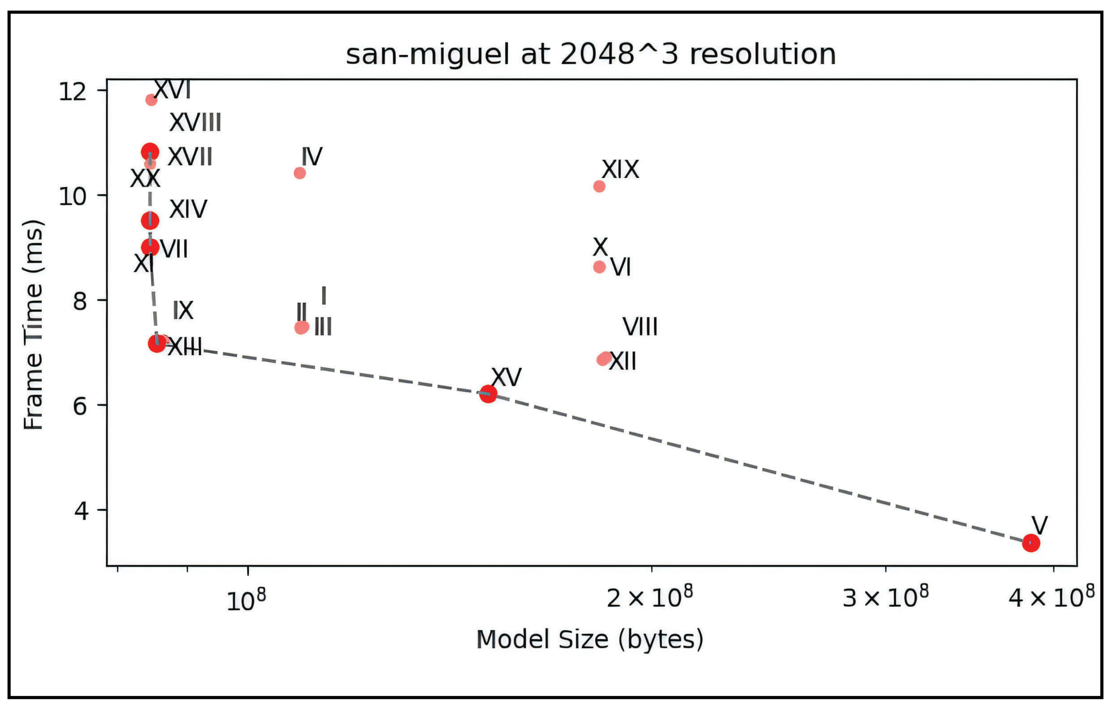

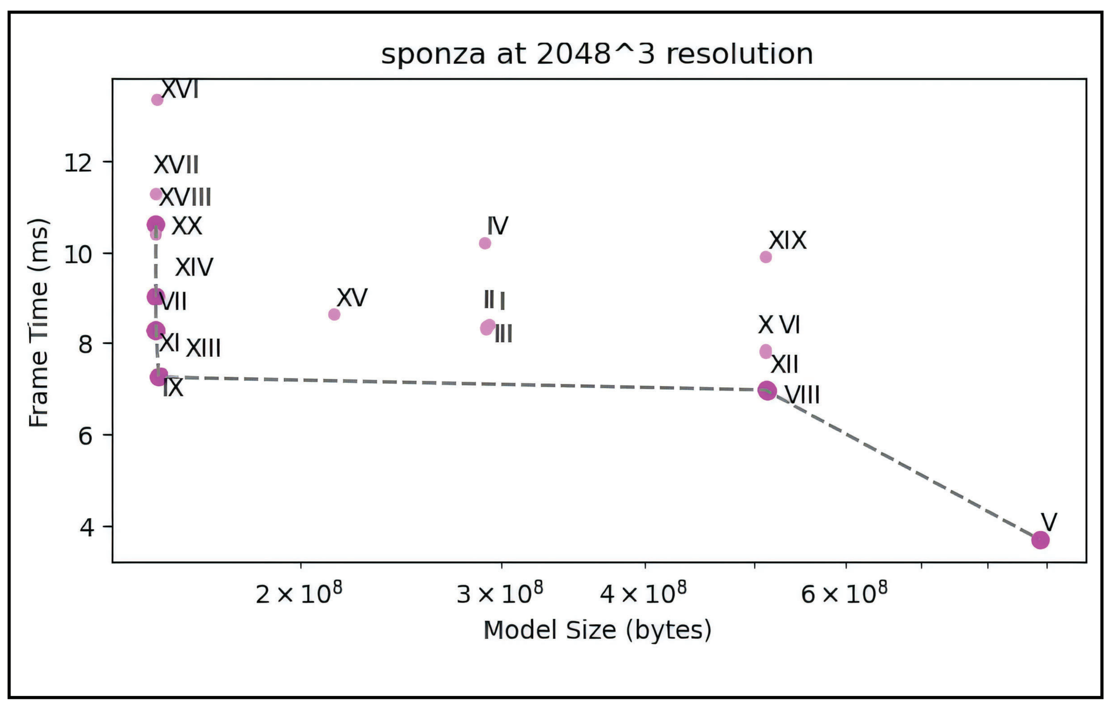

Hybrid formats XI and XIV consistently lie on the Parento frontier of all four models. However, we note that the best format depends on the exact sparsity and homogeneity characteristics of a volume, the resolution, and the desired trade-off between intersection performance and storage.

This motivates future work to systemically explore the search space in an automated fashion. It would also be interesting to investigate how well deep learning models could generate hybrid format proposals for different voxel models. This work describes an infrastructure capable of generating the dataset for training such a model.