Data dependent optimizations for permutation volume rendering

Craig M. Wittenbrink[a] and Kwansik Kim[b]

[a]Hewlett-Packard Laboratories, Palo Alto, CA
[b]Computer Science Dept., University of California, Santa Cruz

Outline

- Main Contributions
- Related work
- Linear octrees
- Performance
- Quality performance tradeoffs
- Conclusions

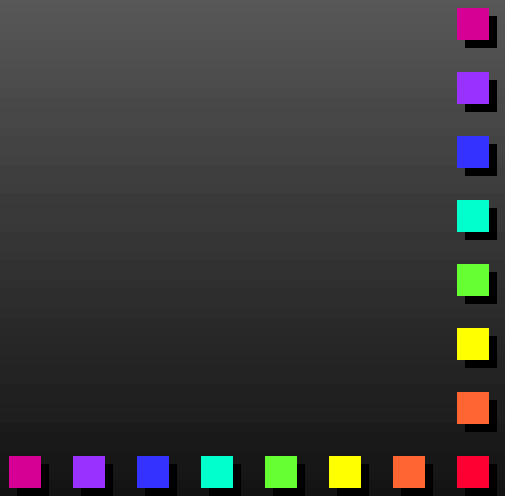

Main contributions

- 3 to 5X Speedup using data dependent acceleration
- 39.3 Mvoxels/second absolute fastest result on MasPar
- Quality time tradeoffs developed

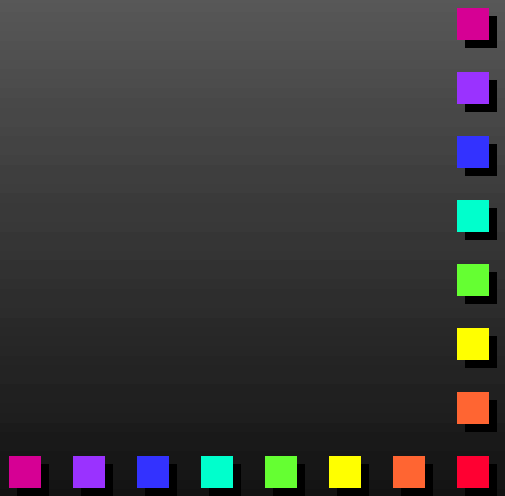

Related work, Mvoxels/second (MasPar MP-1 performance)

- Wittenbrink and Somani, Wz, PRS'93
- Hsu, Hz, PRS'93
- Vezina, Vz, Vol. Vis'92
- Goel*, Visual Computer'96.

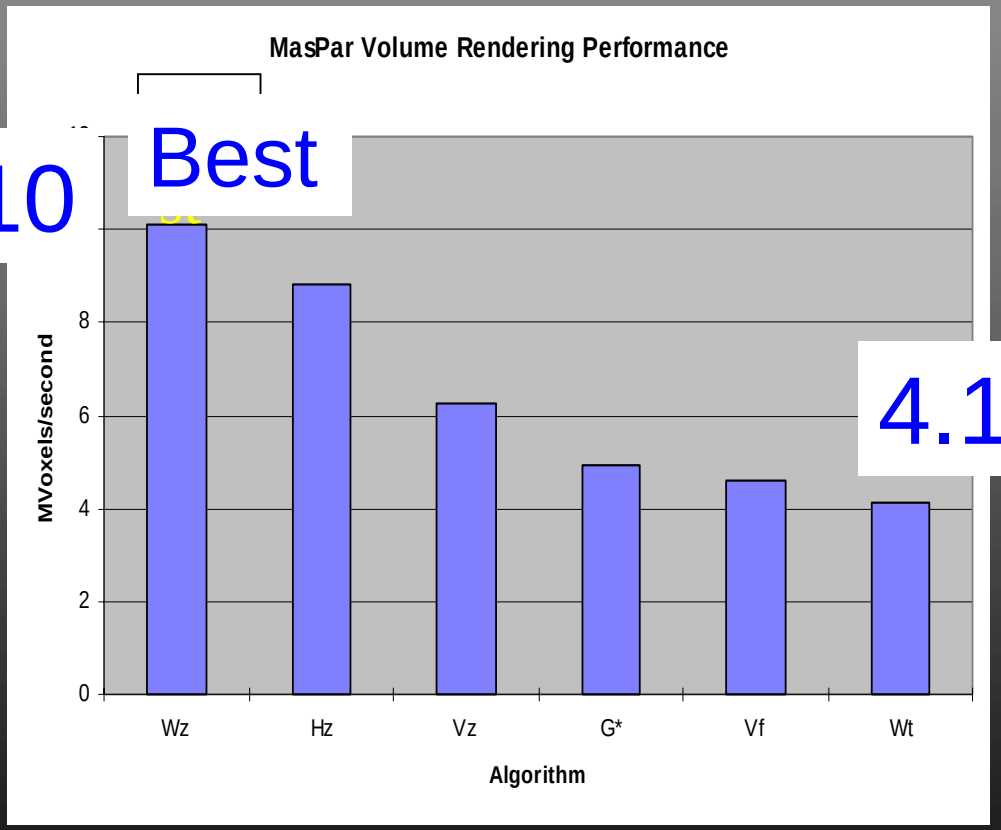

MasPar MP-1 Speedup 16k over 1k processors

- Wz, Supralinear
- Wt, linear
- Vf, Vz, G*, Hz follow

21.3

Ideal 16

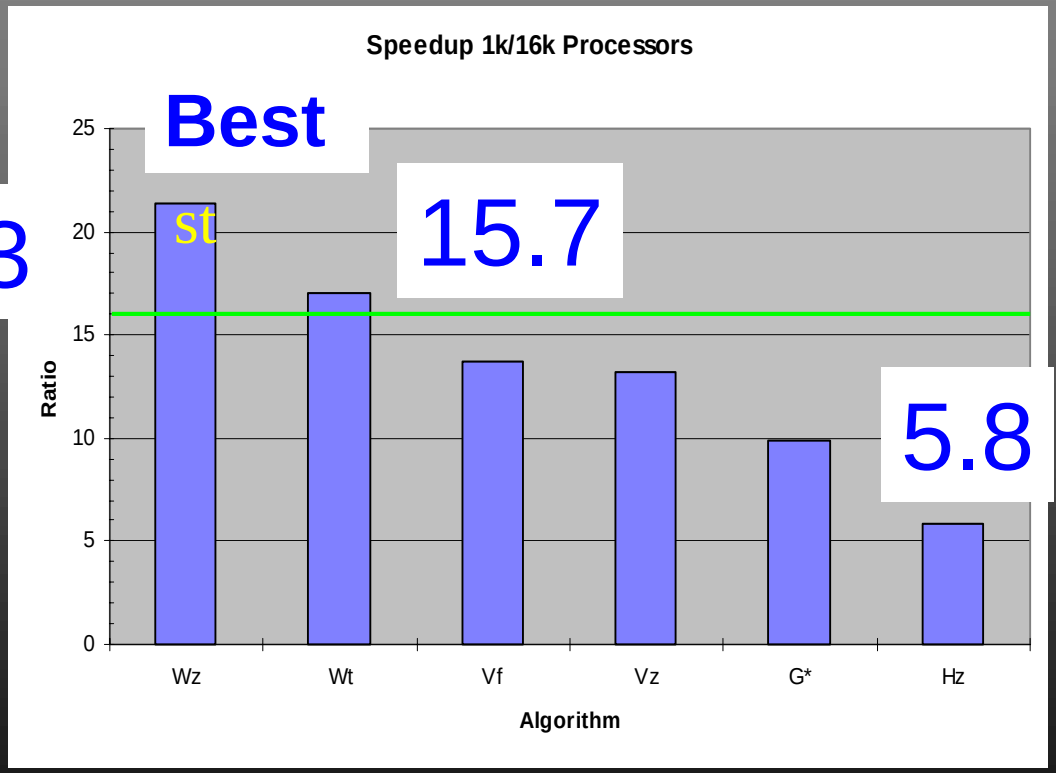

Background, permutation warping

- Regular Grids
- M, one-to-one mapping
- T^{-1} inverse view transform

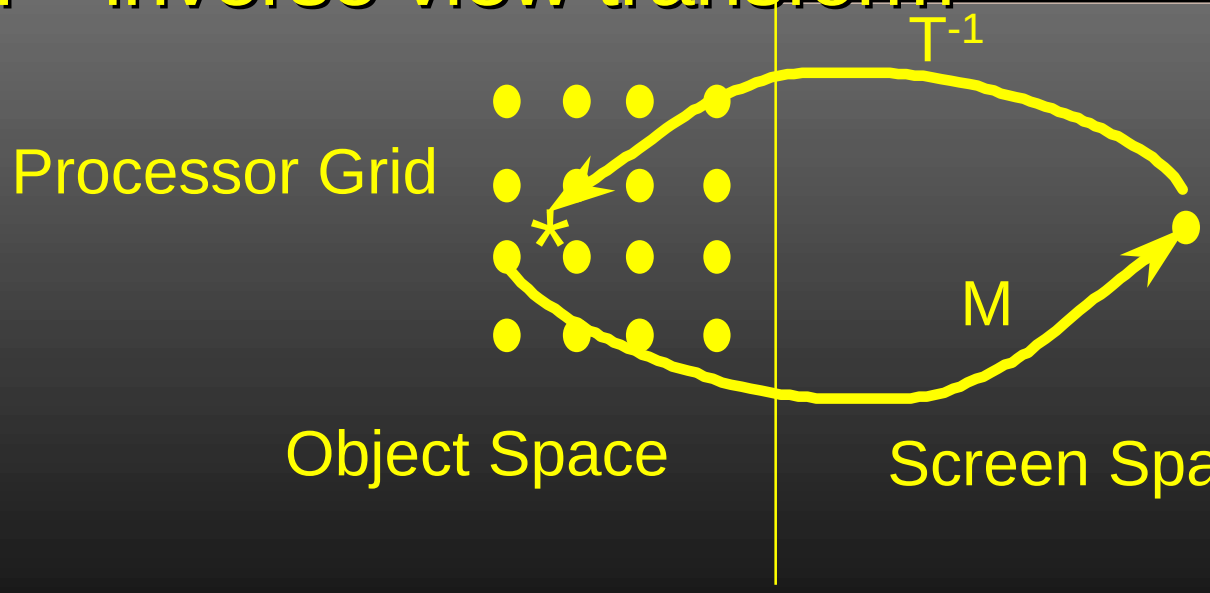

Assignment is always 1 to 1

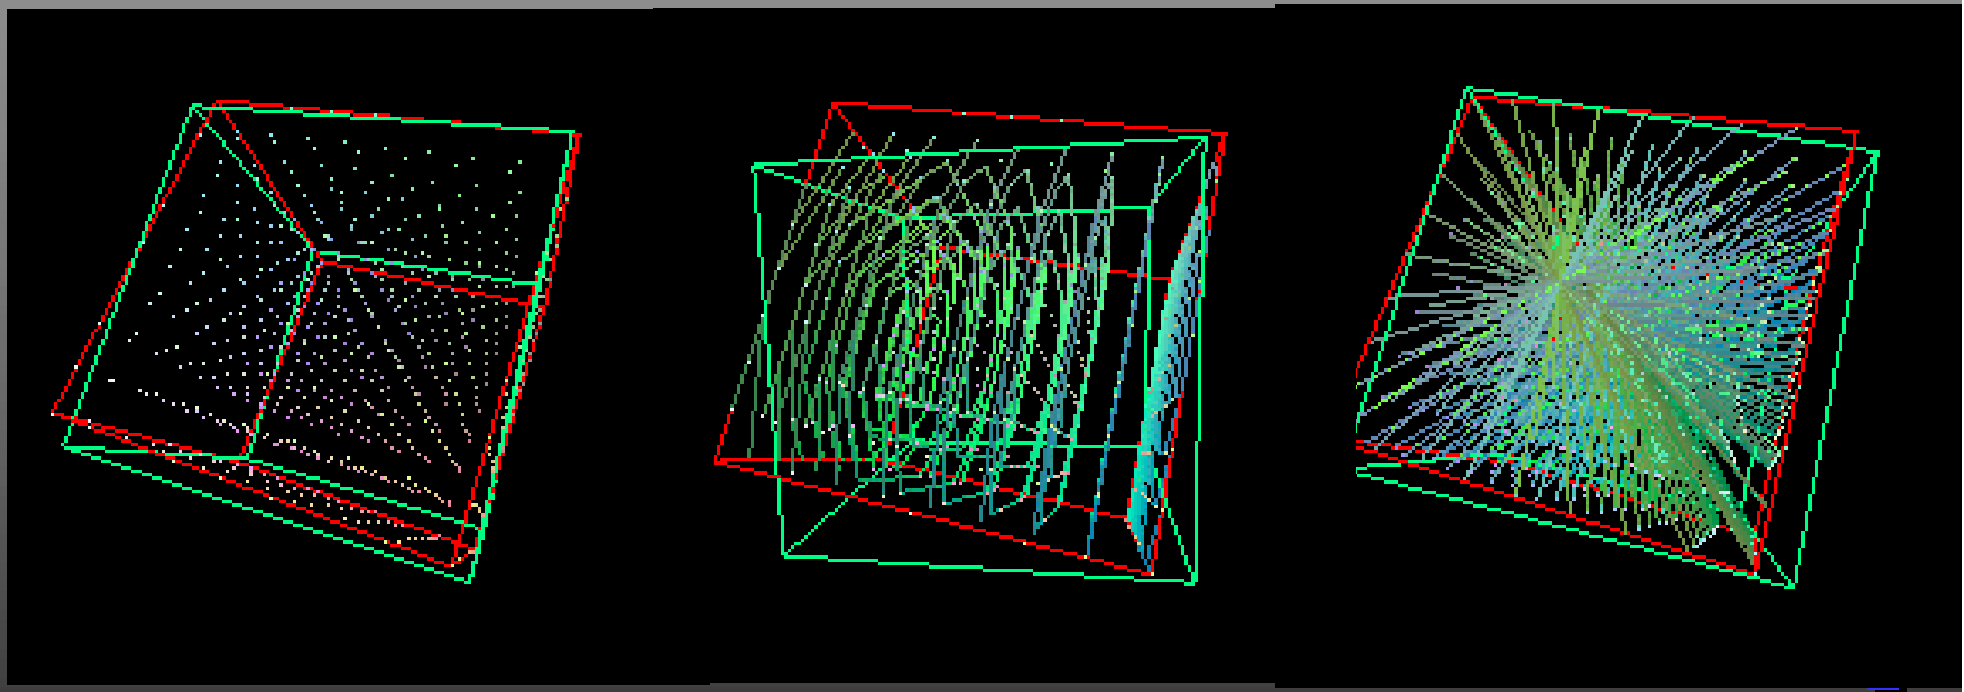

Linear octrees (quadtree example)

4 Linear Quadtrees

PE0: 00,01,02,03,10,11,12,13,20,21,22,23,3X

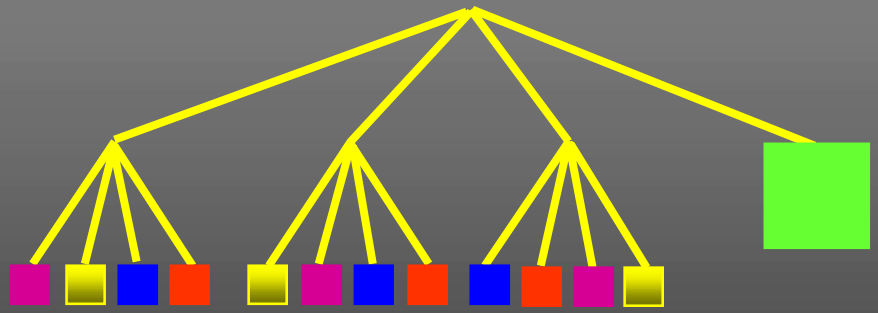

PE3: XX

PE1,PE2: NULL

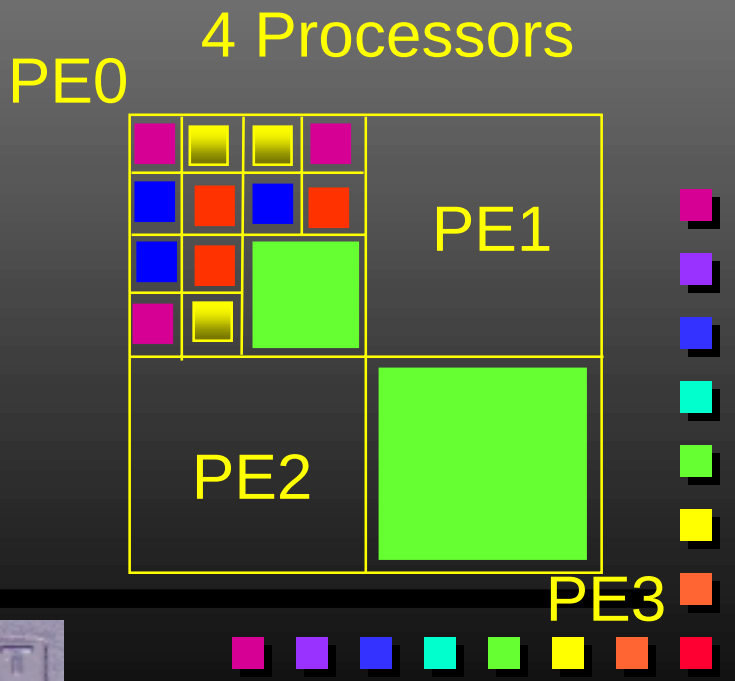

Static load balancing

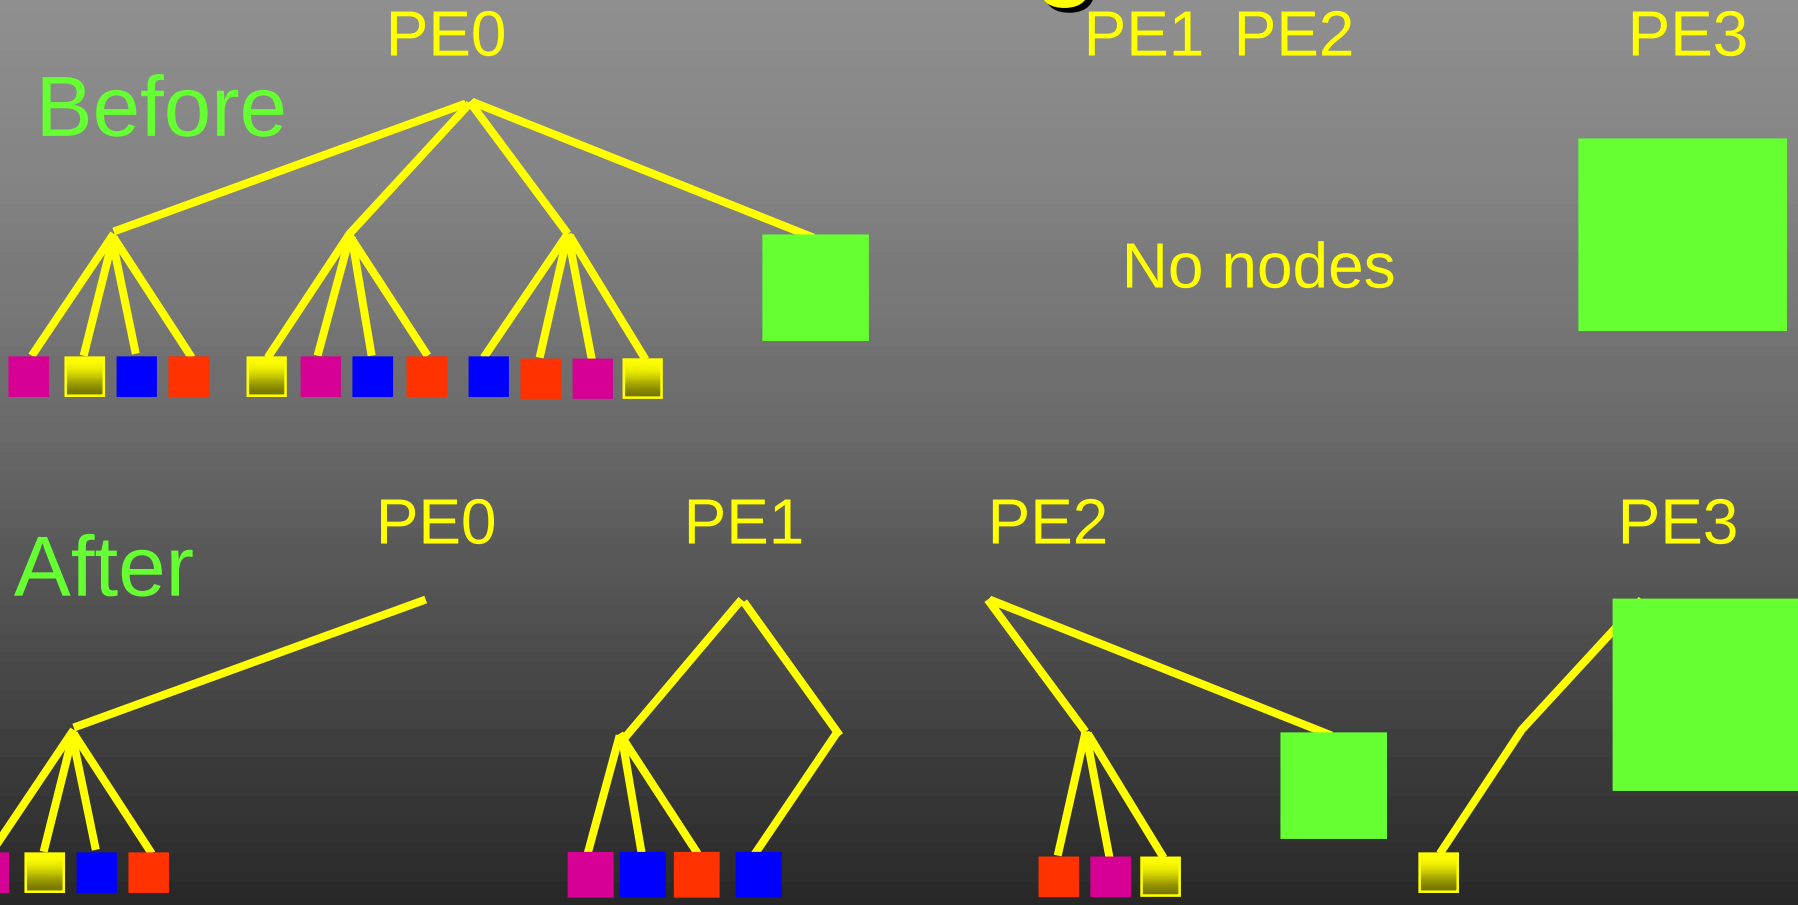

Octree pseudo-code

- Create, condense, and load balance octrees
- Sort nodes by octree level
- while(local nodes) {
  - pick U, U'=MU, $p=T^{-1}$ U'
  - value=interpolateOctree(p)
  - for all 3D voxels in U_i {
    - $U'_j=MU_i$ // permutation assignment
    - send(value,U'_j)}}

Performance achieved MP2 NASA Goddard

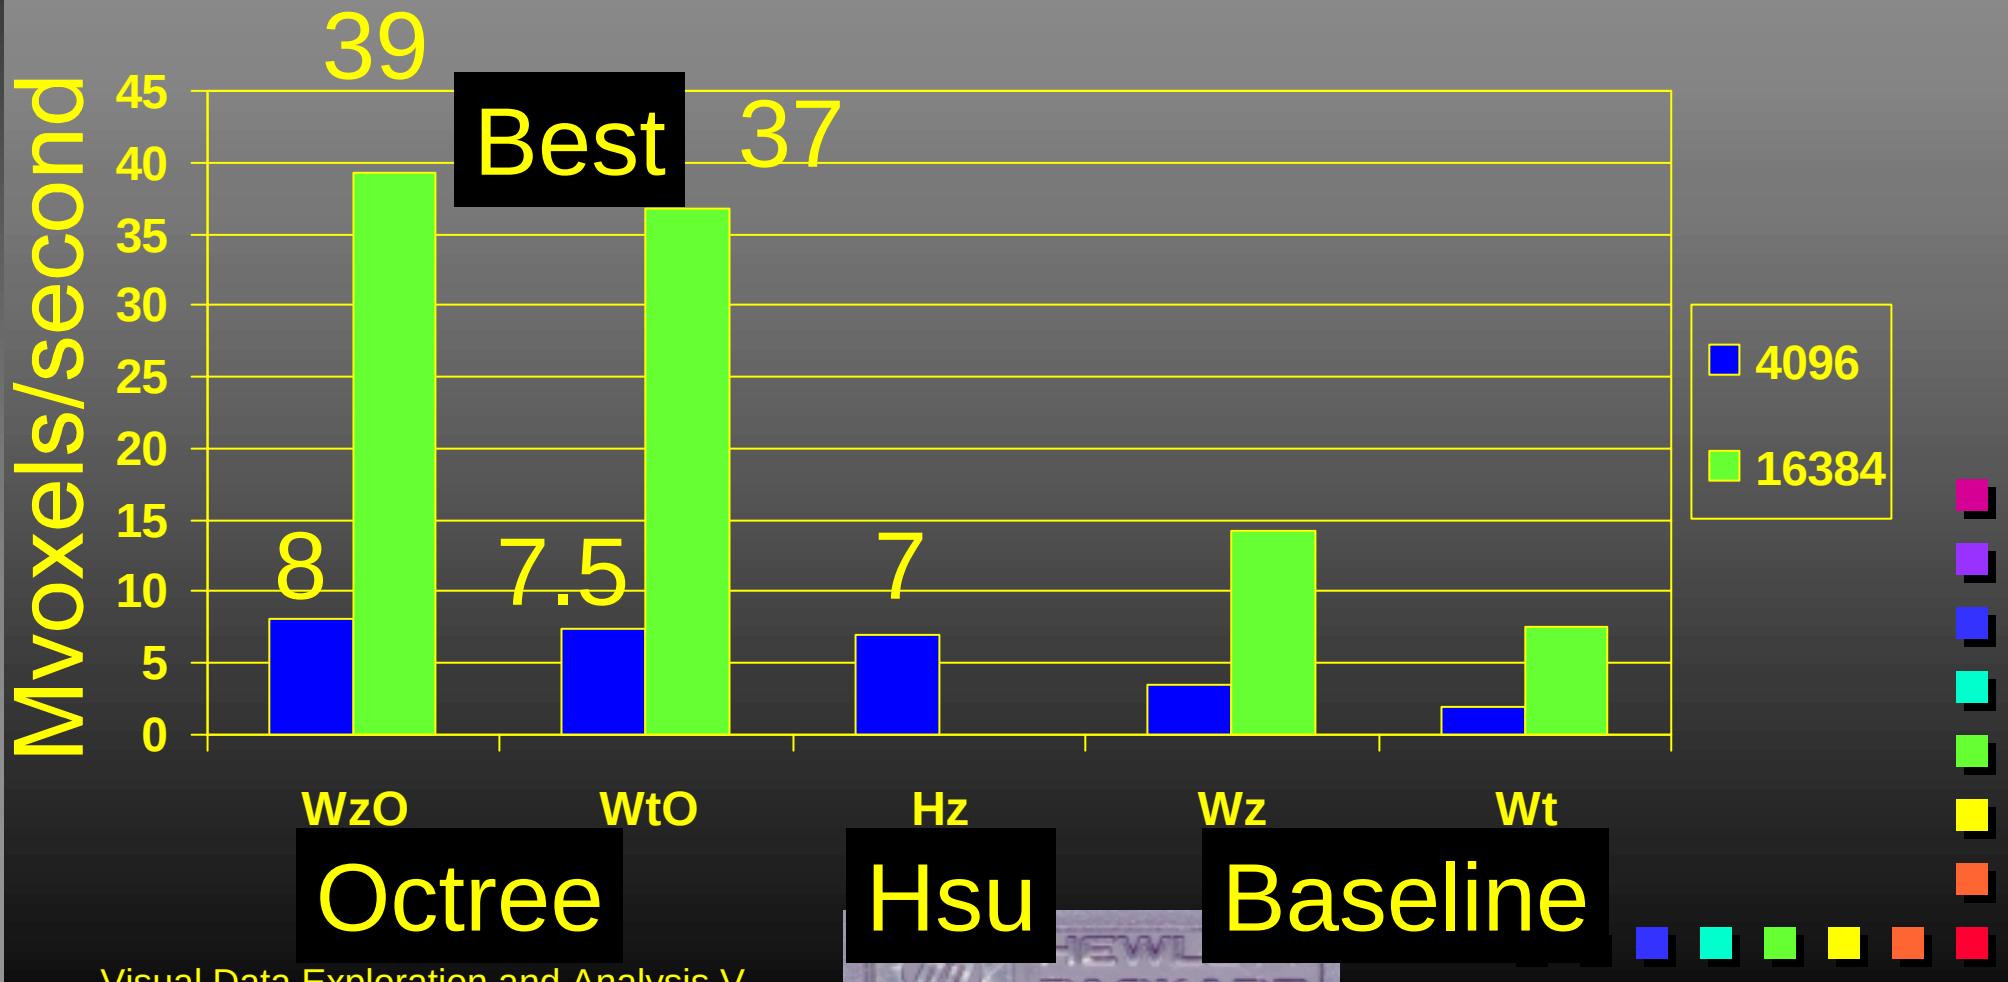

Error versus threshold

time (milliseconds)	1867	1751	804	710	456
theshold	0	1	5	10	50
RMS error	.87	.92	2.7	3.0	7.2

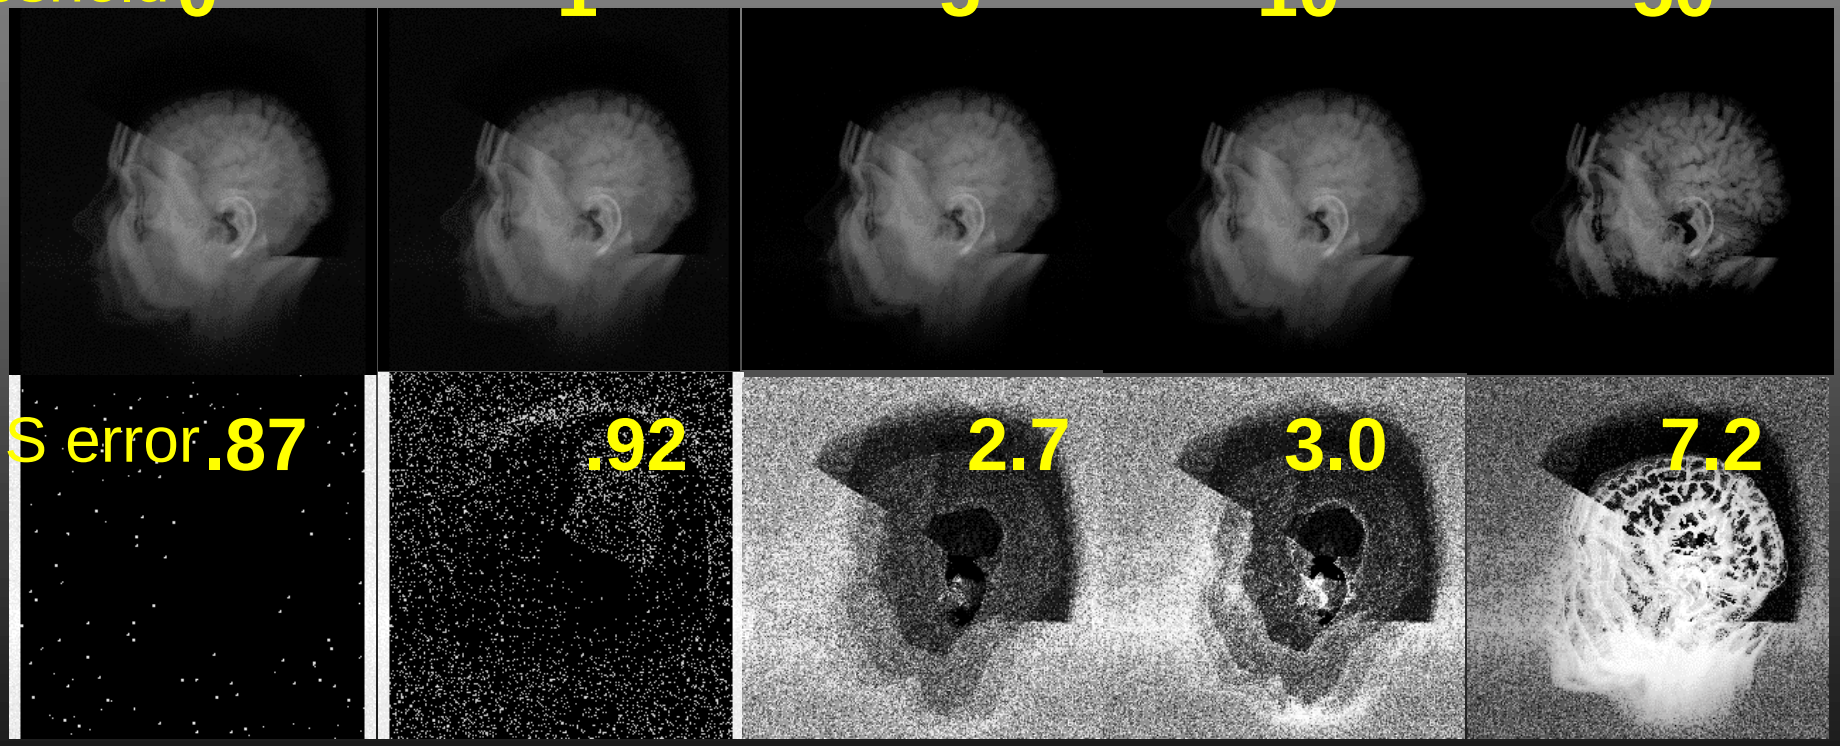

0 and 50 Threshold

0.868 RMS Error

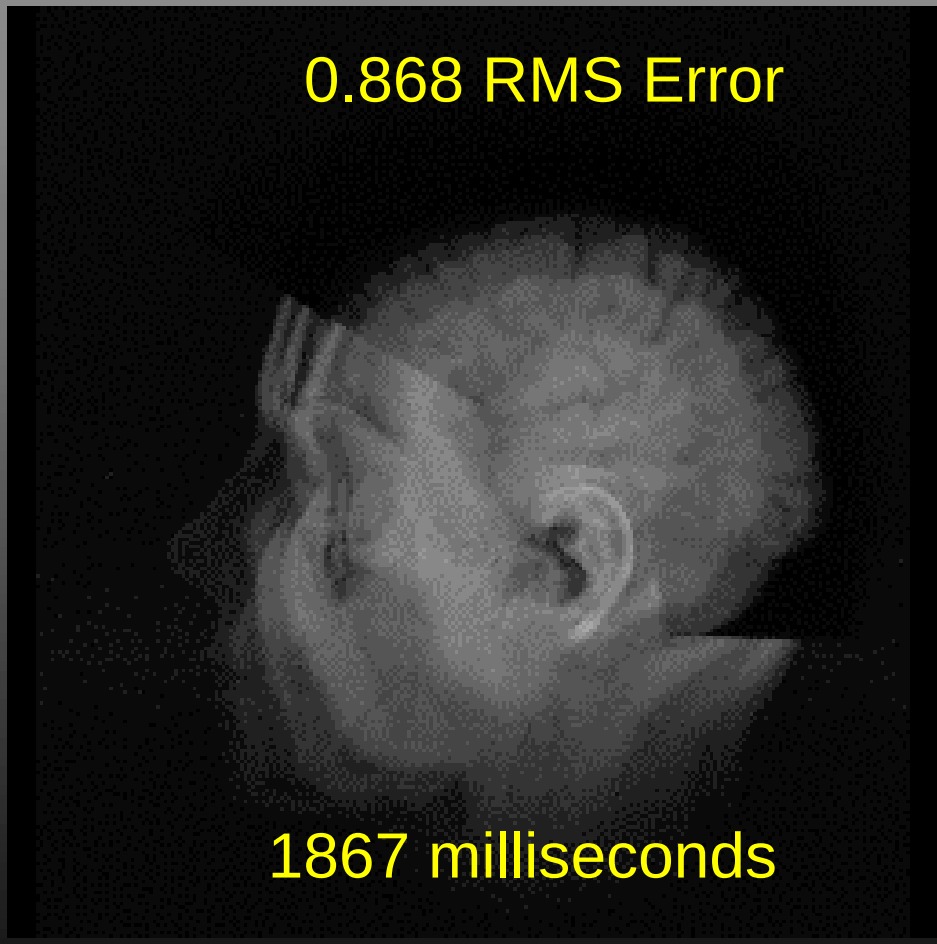

1867 milliseconds

7.218 RMS Error

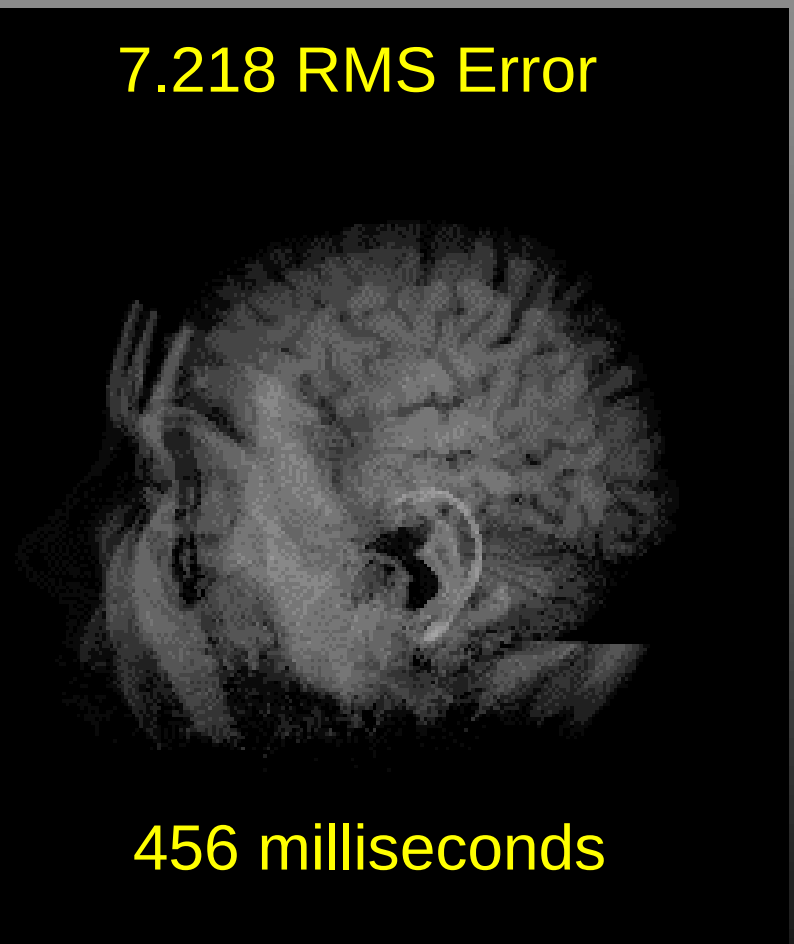

456 milliseconds

Error versus Run time

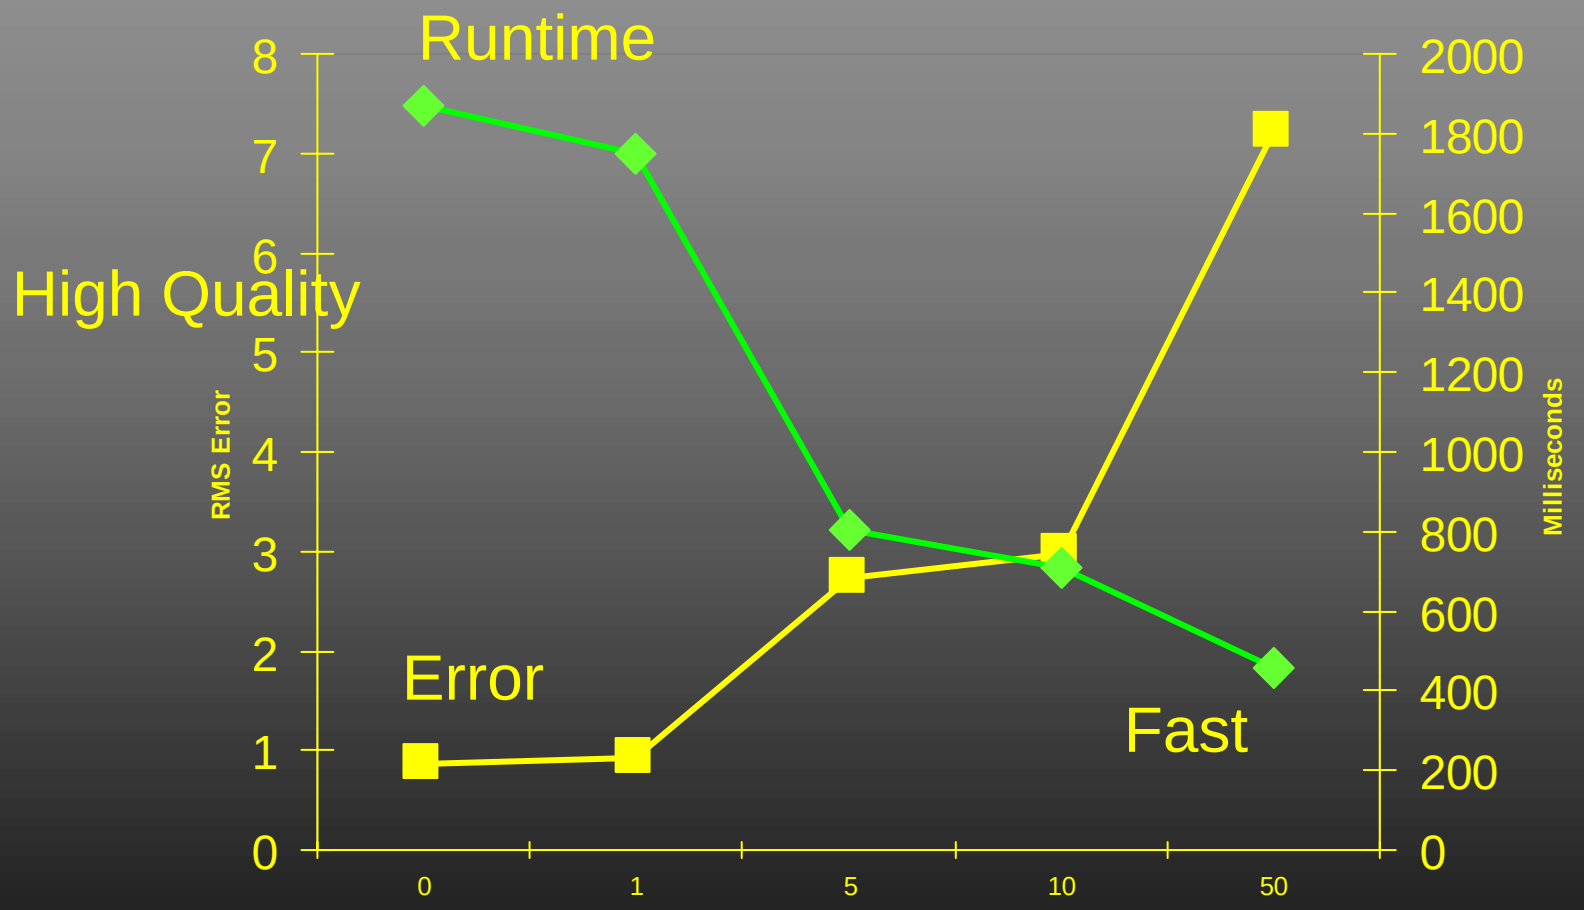

Speedups of 16k over 4k MP-2 Ideal is 4.0 (128^3 volume)

Octree	**3.6**
Baseline tril	**3.85**
Baseline zoh	**4.05**

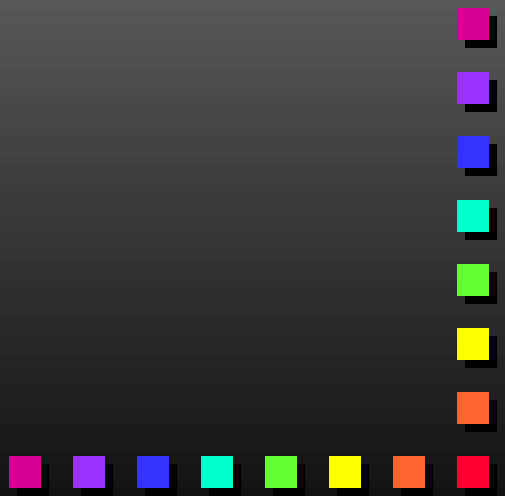

Conclusions

- 3-5X Improvement over baseline
- 39.3 Mvoxels/second fastest MasPar renderer
- One of the highest fidelity (trilinear/higher possible)
- Possibly superior to shear warp for massive parallelism
- Linear octrees effective on SIMD