

HLBVH: Hierarchical LBVH Construction for Real Time Ray Tracing of Dynamic Geometry

Jacopo Pantaleoni and David Luebke
NVIDIA Research

- **Real Time Ray Tracing is almost there***

[Garanzha and Loop 2010, Aila and Laine 2009, Wald et al 2007, ...]

160-200 M rays/s on GF480



- Real Time Ray Tracing is almost there*

[Garanzha and Loop 2010, Aila and Laine 2009, Wald et al 2007, ...]

160-200 M rays/s on GF480

* but only for **static** scenes



- Real Time Ray Tracing is almost there*

[Garanzha and Loop 2010, Aila and Laine 2009, Wald et al 2007, ...]

160-200 M rays/s on GF480

* but only for **static** scenes



- Spatial Index construction real-time only for **100K** tris!

- Real Time Ray Tracing is almost there*

[Garanzha and Loop 2010, Aila and Laine 2009, Wald et al 2007, ...]




160-200 M rays/s on GF480

* but only for **static** scenes



- Spatial Index construction real-time only for **100K** tris!

- Our target is **1M** dynamic tris




- **Many approaches: refitting, partial rebuilds...
but LBVH [Lauterbach et al] probably fastest
available GPU builder**


- Many approaches: refitting, partial rebuilds...
but LBVH [Lauterbach et al] probably fastest
available GPU builder
- still not fast enough... 1M tris => **~150ms**

- Many approaches: refitting, partial rebuilds...
but LBVH [Lauterbach et al] probably fastest
available GPU builder
- still not fast enough... 1M tris => ~**150ms**
- But could be made faster! 😊



- Consider barycenters of each primitive






- Consider barycenters of each primitive so that it works with point sets



- Consider barycenters of each primitive so that it works with point sets
- sort them along a 1D Morton curve through a grid...



- Consider barycenters of each primitive so that it works with point sets
- sort them along a 1D Morton curve through a grid...
- and group them by cell




- Morton codes computed using 10 bits per component
- primitives sorted with a single 30bit **global** sort
- parallel hierarchy emission required 2 additional sorting operations on $\Omega(N * 30)$ split planes

H(ierarchical)LBVH

HLBVH: at a glance

- hierarchical process



- exploit spatial and temporal coherence in the input mesh

- novel hierarchy emission algorithm

- novel SAH hybrid



HLBVH: primitive sorting

- Given a point its *Morton code* is obtained interleaving the bits of its coordinates:

e.g. (0**1**00, 1**0**01, 0**1**11) => 010**1**0100**1**011

- Each triplet of bits => next octant in a grid hierarchy:


2D example: 01**11**



- Consider a 2 level hierarchy:

coarse: $3m$ bits


fine: $3n$ bits



- Consider a 2 level hierarchy:

coarse: $3m$ bits

fine: $3n$ bits




- smaller m \Rightarrow higher chances consecutive prims fall in the same voxel (e.g. $\{1,2\}$, $\{3,4\}$)

- Consider a 2 level hierarchy:

coarse: $3m$ bits

fine: $3n$ bits




- smaller m => higher chances consecutive prims
fall in the same voxel (e.g. {1,2}, {3,4})

- Exploit coherence:

Compress-Sort-Decompress [Garanzha and Loop 2010]


within **coarse** grid

- Compute **n**-bit Morton codes




- Compress: run-length encode based on first **3m** bits

- Compute **n**-bit Morton codes



- Compress: run-length encode based on first **3m** bits
- Sort: do a **3m**-bit radixsort of the rle key blocks


- Compute n -bit Morton codes



- Compress: run-length encode based on first $3m$ bits
- Sort: do a $3m$ -bit radixsort of the rle key blocks
- Decompress: run-length decode sorted keys

- CSD at work:

$\{ 7, 7, 1, 1, 1, 3, 3, 4, 5, 5 \}$



HLBVH: primitive sorting (part 1)


- CSD at work:

$\{ 7, 7, 1, 1, 1, 3, 3, 4, 5, 5 \}$

- Compress:

$\{ 7, 1, 3, 4, 5 \}$ run values

$\{ 2, 3, 2, 1, 2 \}$ run lengths



- CSD at work:

$\{ 7, 7, 1, 1, 1, 3, 3, 4, 5, 5 \}$

- Compress:


$\{ 7, 1, 3, 4, 5 \}$ run values

$\{ 2, 3, 2, 1, 2 \}$ run lengths

- Sort:

$\{ 1, 3, 4, 5, 7 \}$ run values

$\{ 3, 2, 1, 2, 2 \}$ run lengths




- CSD at work:

$\{ 7, 7, 1, 1, 1, 3, 3, 4, 5, 5 \}$

- Compress:

$\{ 7, 1, 3, 4, 5 \}$ run values

$\{ 2, 3, 2, 1, 2 \}$ run lengths



- Sort:


$\{ 1, 3, 4, 5, 7 \}$ run values

$\{ 3, 2, 1, 2, 2 \}$ run lengths


- Decompress:

$\{ 1, 1, 1, 3, 3, 4, 5, 5, 7, 7 \}$


- Meshes often show such coherence




- Prims are now sorted
in **coarse** voxels




- Prims are now sorted in **coarse** voxels



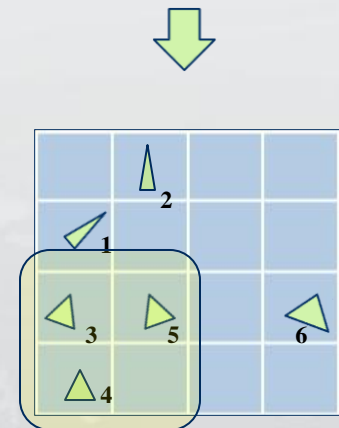
- Sort within each voxel using intra-cta (shared-mem) sort



- By CSD we have substantially reduced **BW** taking advantage of **spatial coherence**



- And if we reuse the final ordering across frames, we can take advantage of **temporal coherence** too



HLBVH: hierarchy emission

- This is all good, but we are still left with hierarchy emission, which is the hard part:

hierarchy emission

$2 * \Omega(N*30)$ sorts

vs

prim sorting

$1 * O(N)$ sort

in LBVH

- **Input: array of sorted prims**
- **Output: array of nodes forming a tree**

- Input: array of sorted prims
(sequence of Morton codes)

0	1	2	3	4	...	N-4	N-3	N-2	N-1
0	0	0	0	0	...	1	1	1	1
0	0	0	0	0	...	0	0	0	1
0	0	0	1	0	...	0	0	1	1
0	0	0	1	1	...	1	1	1	0
1	1	1	0	0	...	1	1	1	0
0	0	1	0	0	...	0	0	0	0
1	1	0	0	0	...	0	1	0	0
0	1	0	1	1	...	0	0	0	0

- Input: array of sorted prims
(sequence of Morton codes)
- Output: sequence of nested *segments*


[0]	1	2	3	4	...	[N-4]	N-3	N-2	N-1]
0	0	0	0	0	...	1	1	1	1
0	0	0	0	0	...	0	0	0	1
0	0	0	1	0	...	0	0	1	1
0	0	0	1	1	...	1	1	1	0
1	1	1	0	0	...	1	1	1	0
0	0	1	0	0	...	0	0	0	0
1	1	0	0	0	...	0	1	0	0
0	1	0	1	1	...	0	0	0	0

- Input: array of sorted prims
(sequence of Morton codes)
- Output: sequence of nested *segments*

[0]	1	2]	[3	4	...	[N-4	N-3]	[N-2	N-1]
0	0	0	0	0	...	1	1	1	1
0	0	0	0	0	...	0	0	0	1
0	0	0	1	0	...	0	0	1	1
0	0	0	1	1	...	1	1	1	0
1	1	1	0	0	...	1	1	1	0
0	0	1	0	0	...	0	0	0	0
1	1	0	0	0	...	0	1	0	0
0	1	0	1	1	...	0	0	0	0

- Partial Breadth First Traversal

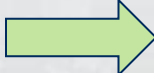
- Consider **p**-bit planes
at a time



0	1	2	3	4	...	N-4	N-3	N-2	N-1
0	0	0	0	0	...	1	1	1	1
0	0	0	0	0	...	0	0	0	1
0	0	0	1	0	...	0	0	1	1
0	0	0	1	1	...	1	1	1	0
1	1	1	0	0	...	1	1	1	0
0	0	1	0	0	...	0	0	0	0
1	1	0	0	0	...	0	1	0	0
0	1	0	1	1	...	0	0	0	0

- Partial Breadth First Traversal

- Consider **p**-bit planes
at a time




0	1	2	3	4	...	N-4	N-3	N-2	N-1
0	0	0	0	0	...	1	1	1	1
0	0	0	0	0	...	0	0	0	1
0	0	0	1	0	...	0	0	1	1
0	0	0	1	1	...	1	1	1	0
1	1	1	0	0	...	1	1	1	0
0	0	1	0	0	...	0	0	0	0
1	1	0	0	0	...	0	1	0	0
0	1	0	1	1	...	0	0	0	0

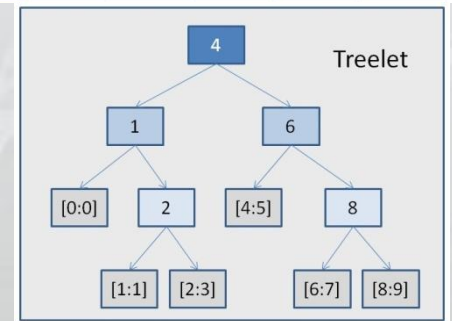
- Partial Breadth First Traversal

- Consider **p**-bit planes
at a time

0	1	2	3	4	...	N-4	N-3	N-2	N-1
0	0	0	0	0	...	1	1	1	1
0	0	0	0	0	...	0	0	0	1
0	0	0	1	0	...	0	0	1	1
0	0	0	1	1	...	1	1	1	0
1	1	1	0	0	...	1	1	1	0
0	0	1	0	0	...	0	0	0	0
1	1	0	0	0	...	0	1	0	0
0	1	0	1	1	...	0	0	0	0

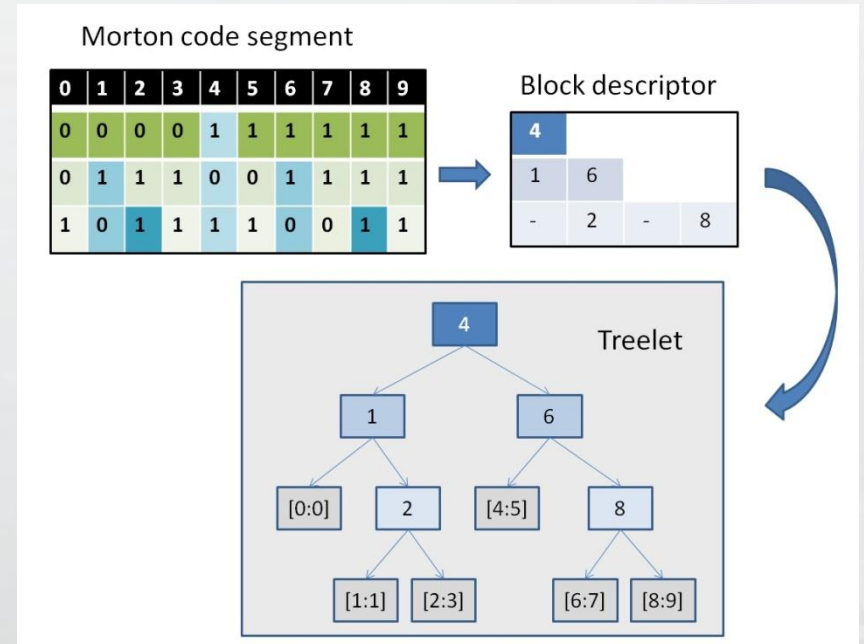


- For each **segment**, emit a **treelet**



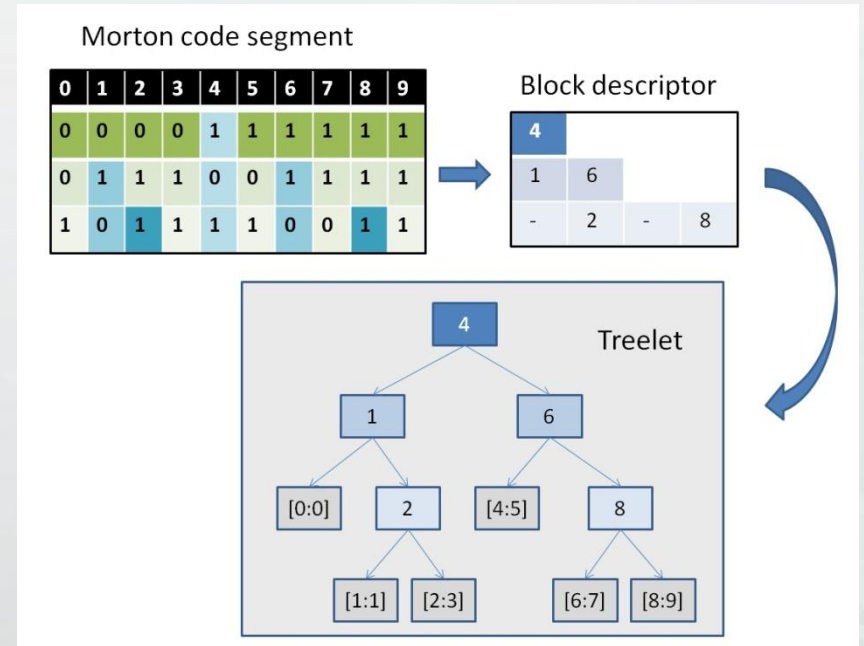
- Partial Breadth First Traversal

- Details in the paper

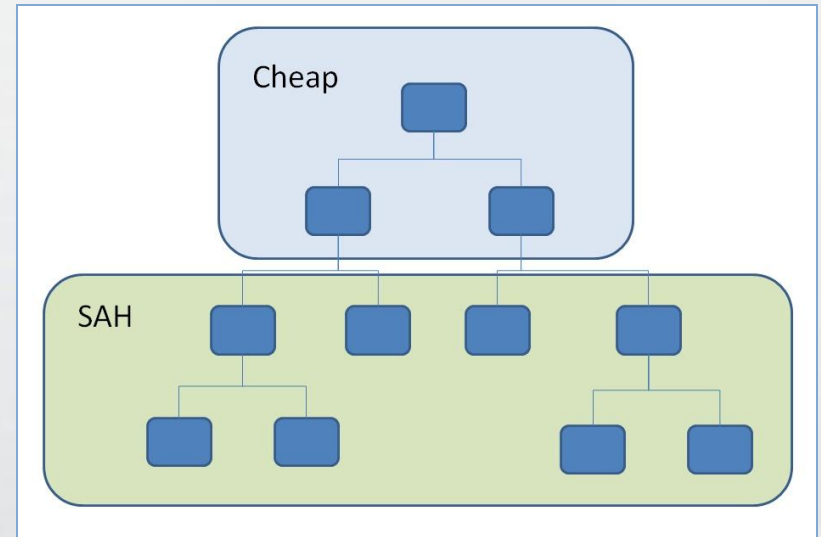


- Partial Breadth First Traversal

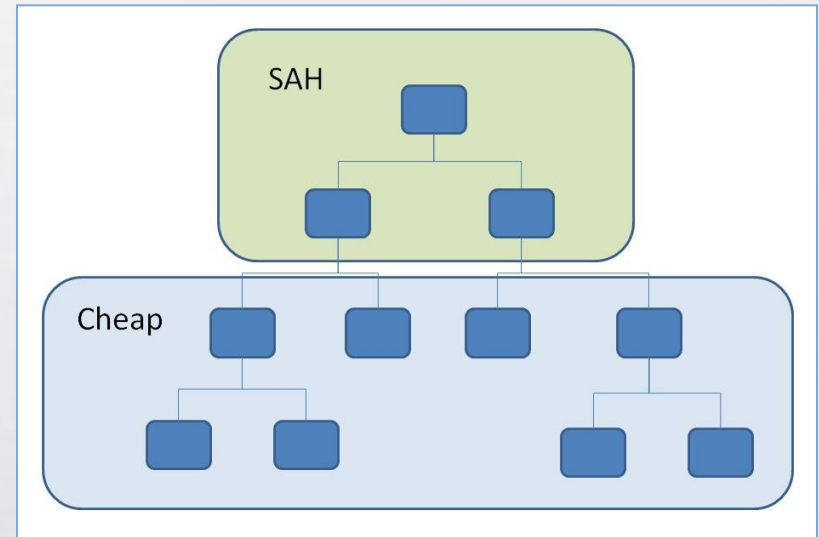
- Details in the paper



- Lauterbach and Wald suggested to perform SAH at the bottom of the tree




- Lauterbach and Wald suggested to perform SAH at the bottom of the tree
- But with CSD we can do better!
Our coarse clusters can be used to build a SAH-based top-level tree



- Lauterbach and Wald suggested to perform SAH at the bottom of the tree

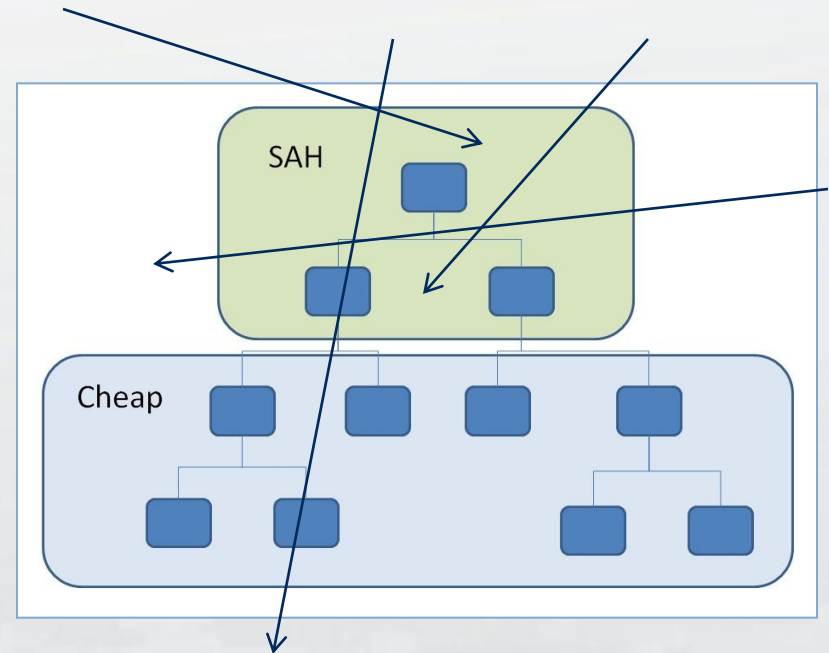
- But with CSD we can do better!

Our coarse clusters can be used to build a SAH-based top-level tree




- As the clusters are few, the overhead is **low**

- Not only this is faster...
- It's also better because the top-level tree is what matters mostly



- We reduced **BW** by **>10x**
- We exploit spatial and temporal coherence
- Support fully dynamic geometry, from deformations to chaotic fracturing
- Low-overhead SAH hybrid


- 1M fully dynamic tris => **~35ms**



- 2M incoherent



- 350k coherent



- Cleanly coded using Thrust

- Will be available at:

<http://code.google.com/p/hlbvh/>

```
hierarchy_emission(codes, N_prims, n_bits)
1  int segment_heads[]
2  int head_to_node[ N_prims ] = {-1}
3  head_to_node[0] = segment_heads[0] = 0
4
5  for (level = 0; level < n_bits; level += p)
6  // compute segment ids
7  segment_id[i] = scan( head_to_node[i] ≠ -1)
8
9  // get the number of segments
10 int N_segments = segment_id[N_prims-1]
11
12 int P = (1 << p) - 1
13
14 // compute block descriptors
15 int block_splits[ N_segments * P ] = {-1}
16 foreach i in [0, N_prims)
17   emit_block_splits(
18     i, [in] primitive index to process
19     codes, [in] primitive Morton codes
20     [level, level + p), [in] bit planes to process
21     segment_id, [in] segment ids
22     head_to_node, [in] head to node map
23     segment_heads, [in] segment heads
24     block_splits ) [out] block descriptors
25
26 // compute the block offsets summing
27 // the number of splits in each block
28 int block_offsets[ N_segments + 1 ]
29 block_offsets[s] = ex_scan( count_splits(s) )
30 int N_splits = block_offsets[N_segments]
31
32 // emit treelets and update
33 // segment_heads and head_to_node
34 foreach segment in [0, N_segments)
35   emit_treelets(
36     segment, [in] block to process
37     block_splits, [in] block descriptors
38     block_offsets, [in] block offsets
39     segment_id, [in] segment ids
40     head_to_node, [in/out] head to node map
41     segment_heads ) [in/out] segment heads
42
43 node_count += N_splits * 2
```

Figure 4: Pseudocode for our hierarchy emission loop.

Thank You!