GPU Accelerated Tandem Traversal of Blocked Bounding Volume Hierarchies

Jesper Damkjær and Kenny Erleben
{damkjaer,kenny}@diku.dk

Department of Computer Science
University of Copenhagen

October 2009

Traditional BVH Traversal

- Two BVHs are traversed
    - Using either a stack or a queue
    - Using a descend rule descending either tree
    - Descend both trees simultainiously
- For each descend, the BVs in the nodes are compared for overlap

Naive BVH on GPU

- One pair of BVHs per Thread
- Upper space bound for stack

$$k\,(c-1)\max\left(\mathbf{height}(A),\mathbf{height}(B)\right),$$

max. cardinality, c, and size of two BV node references, k.

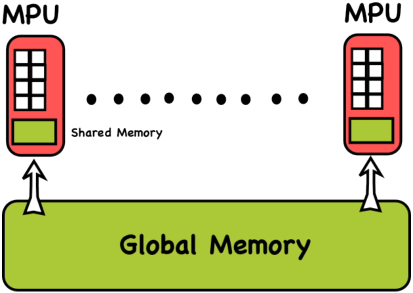

Shared memory too small and global memory too slow

Use Blocks

- 1 Block $\equiv$ Each node has 4 children
- If overlap $\Rightarrow$ 16 new overlaps

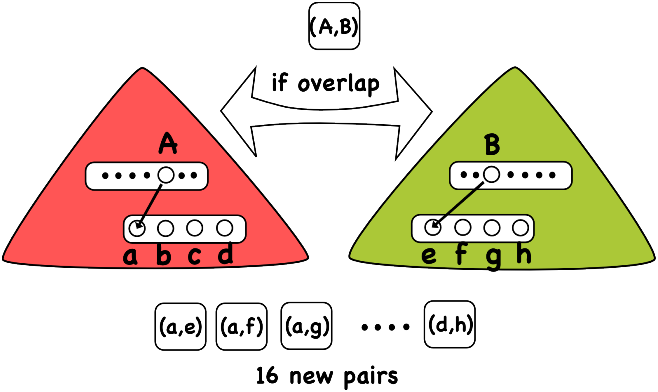

Less data to transfer and more work per thread

Use Double Buffered List

Stack/Queue ⇒ Double buffered list

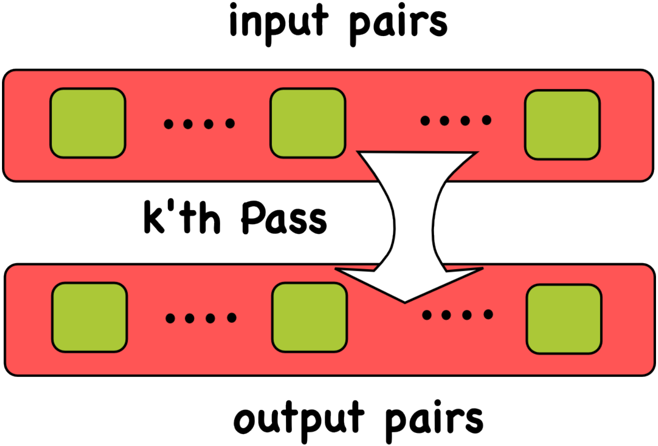

Swap input/output paris for next pass

Memory Trick Needed

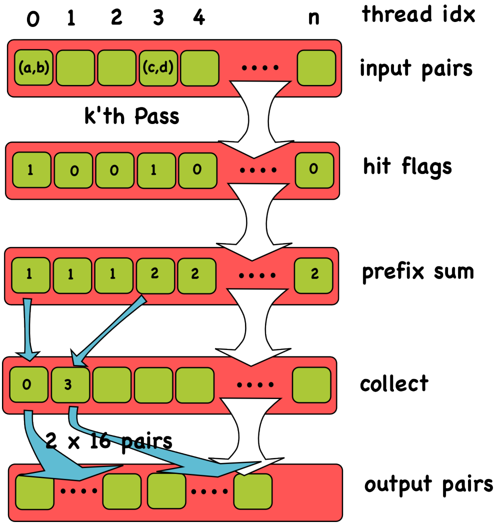

Need Imaginary Nodes

Less than 4 children ⇒ fill with imaginary nodes

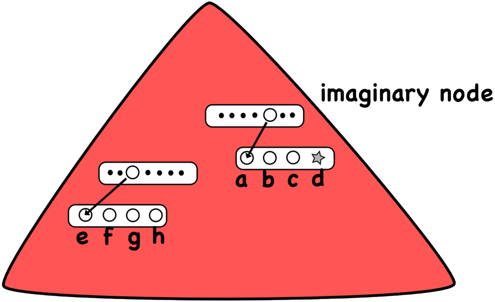

Fills up space ⇒ part of calculation time ⇒ use sparesly

Blocks with Mixed Internal or Leaf Nodes

Not allowed ⇒ Simpler code

Internal Block versus Leaf Block

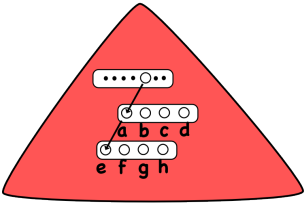

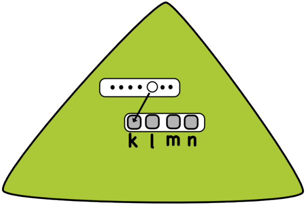

if collide (a, k) ⇒ push (e, k)
if collide (a, l) collision ⇒ push (e, k)
if collide (a, m) collision ⇒ push (e, k)
if collide (a, n) collision ⇒ push (e, k)

Redundant results ⇒ add extra check to code

The Test Setup

Three different configuration types

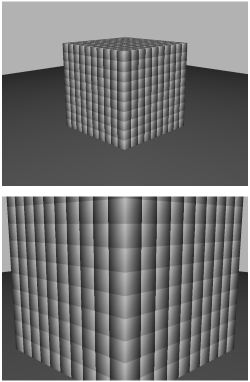

Structured stack

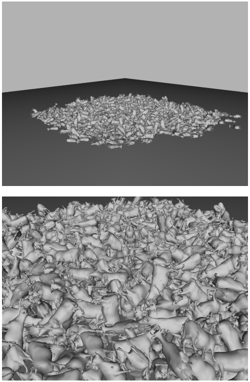

Unstructured Pile

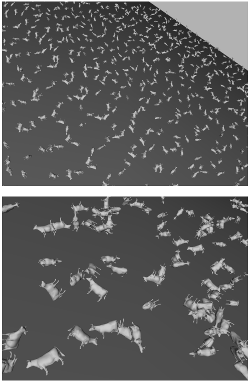

Rock Slide

The Test Setup (Cont'd)

- For each configuration type
  - Increasing number of triangles in objects
  - Increasing number of objects
- Test against Rapid
  - Rapid uses OBBs we use AABBs
- No optimization of imaginary nodes in BVHs (upto 33%)

Results

Rapid on Intel Quad CPU using one core

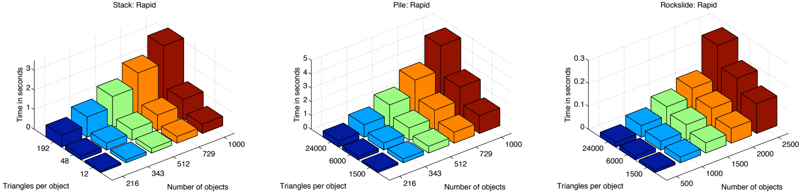
Stack: Rapid

Pile: Rapid

Rockslide: Rapid

Cuda on ge9800 GX2 using one core

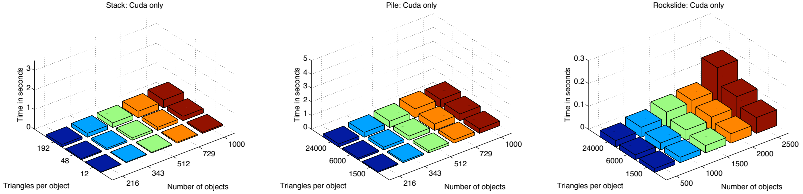
Stack: Cuda only

Pile: Cuda only

Rockslide: Cuda only

Stack (5-8) Pile (3-7) Slide (2)

Thanks

Questions?