

DAG Construction

Erik Sintorn

Chalmers University of Technology, Sweden

EG EUROGRAPHICS
2018



Welcome back

Before the break, we went through the theory.

These next two talks will be on more practical things.

I will start by talking about how we actually build these DAGs on GPUs, using parallel primitives.

Outline

- Content
 - Voxelization and building an SVO
 - Efficient SVO to DAG compression, first version
 - Efficient SVO to DAG compression, second version

EE

So, the outline of the talk is that:

I will first talk briefly about voxelization

And then about how to create an SVO from the voxelization

Then I will explain how we *used to* build a DAG from the SVO

And then I will go through an alternative way of doing this, which is quite a lot simpler and faster.

Parallel primitives in CUDA

- Our GPU Implementation is based on standard parallel primitives
 - Stream Compaction
 - Sorting
 - Prefix Sum
- A number of libraries exist for doing these things
 - THRUST (<http://docs.nvidia.com/cuda/thrust/index.html>)
 - CUDPP (<http://cudpp.github.io/>)
 - CUB (<https://nvlabs.github.io/cub/>)

EE

Before I start though,

I will talk a lot about these parallel primitives.

These are things like: sorting an array, stream compaction and prefix sum.

Stuff that must be fast, and that you do not want to write yourself.

There are a number of libraries available for this:

THRUST actually ships with CUDA, but the last time I tried it it performed very badly,

So I now use CUB.

Either of these libraries should work though.

Voxelization

- Render geometry to voxel-grid sized rendertarget
 - Render triangles from X, Y or Z axis that gives largest area
 - Single pass, choose axis in geometry shader
 - Avoid cracks with [GL_NV_conservative_raster](#)
 - Output long list of fragments as Morton order integers
 - Simple tutorial at <https://developer.nvidia.com/content/basics-gpu-voxelization>

EE

The first step is to voxelize our scene.

We do this in a single renderpass, with a viewport set up to match the resolution of the final DAG.

For each triangle we choose to project it as seen from the X, Y or Z axis, depending on which gives the largest area.

Also, to avoid cracks, we need to use some sort of conservative rasterization.

This is pretty easy these days, with the `GL_NV_conservative_raster` OpenGL extension.

In the fragment shader, we will simply turn the fragments position into the grid coordinates, and output it to a long list as a *Morton code*.

There are many tutorials on voxelization on the internet, and here's a link for one of them.

Build SVO in CUDA

- Given the Morton cords, v_i , of all set voxels
 - Parent coord is $(v_i \gg 3)$
 - Child index is $(v_i \& 7)$

$$v_i = \begin{array}{c} \text{parent coord} \\ x_N y_N z_N \dots x_1 y_1 z_1 x_0 y_0 z_0 \\ \text{child index} \end{array}$$

EE

We output the voxel's position as a Morton code, because that makes it very easy to mask of the three lowest bits and

<click>

Then the upper bits are the coordinates of the leaf node,

And the lower bits are the index of the voxel within its parent.

Build SVO in CUDA (leaf level)

- Sort all voxels (and remove identical)

001010|001 001010|010 001010|011 001011|001 001011|100 001100|100

EE

Now we have a long list of voxels, and we want to build an SVO. We start by sorting the list, and removing all identical nodes. This is stuff that a library will handle efficiently for us.

Build SVO in CUDA (leaf level)

- Sort all voxels (and remove identical)

001010|001 001010|010 001010|011 001011|001 001011|100 010100|100

- Mark first child of all parents

1 0 0 1 0 1

EE

Then, we mark the *first* voxel of each leafnode. We do this by running through all the voxels, and checking the voxel to the left.

If it has the same parent: We write a zero.

If it has a different parent, this is the first child, so we write a one.

Build SVO in CUDA (leaf level)

- Sort all voxels (and remove identical)

001010|001 001010|010 001010|011 001011|001 001011|100 010100|100

- Mark first child of all parents

1 0 0 1 0 1

- Prefix sum gives position of parents

0 0 0 1 1 2

EE

We can then do a *prefix sum* on this list, which will give us, for every voxel, the index of its parent.

Build SVO in CUDA (leaf level)

- Sort all voxels (and remove identical)

001010|001 001010|010 001010|011 001011|001 001011|100 010100|100

- Mark first child of all parents

1 0 0 1 0 1

- Prefix sum gives position of parents

0 0 0 1 1 2

- Build leaf level cords and childmasks

001010 001011 010100
00001110 00010010 00010000

EE

Now we can build the leaf level of our SVO.

<click>

We first build a list of all the leaf level coordinates, that we will use to continue building the tree

We do this by running through all the voxels, finding the position of the leaf node it ended up in, and writing down the parent coord at that location

<click>

Then we build the actual leaf level of the SVO, that is, the childmasks of all the leafs.

We mask out the index of the voxel in the leaf node, look up the position of the leaf node and set the right bit in the corresponding child-mask

Build SVO in CUDA (internal levels)

- Start from coords in last pass

001|010 001|011 010|100

- Mark first child of all parents

1 0 1

- Prefix sum gives position of parents

0 0 1

- Build next level coords childmasks and pointer

001 010
00001100|0 00010000|2

EE

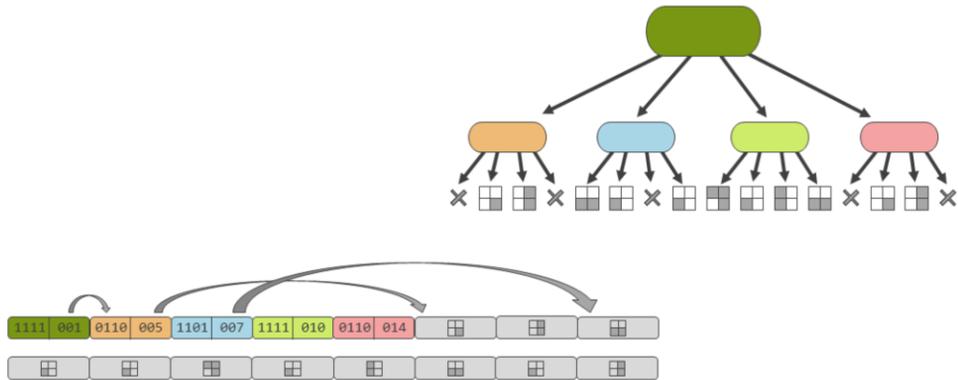
Now we have the leaf level of the SVO, and we need to build all of the internal nodes. We do this in almost exactly the same way.

So, we *start from the leaf coordinates* we created in the last pass,

Mark the first child of all parents

Prefix sum that list to get the position of the parents.

Creating a DAG from an SVO

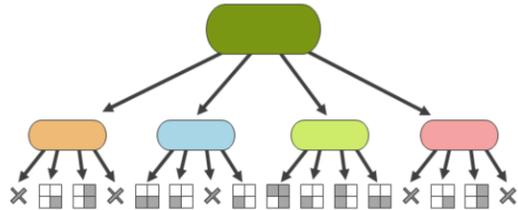


EE

This dag is the 'S'

For brevity, let's say each node takes one word

Creating a DAG from an SVO



Level:

0

1111 001

1

0110 005 1101 007 1111 010 0110 014

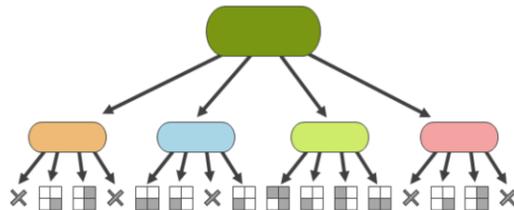
2

12 12 12 12 12 12 12 12 12 12 12 12

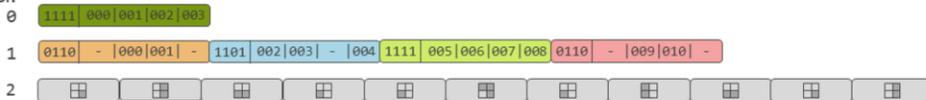
EE

To make this easier, let's visualize the words divided into levels

Creating a DAG from an SVO



Level:



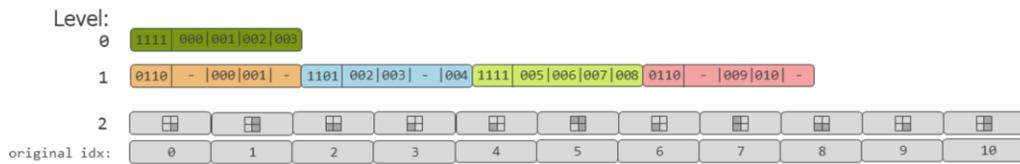
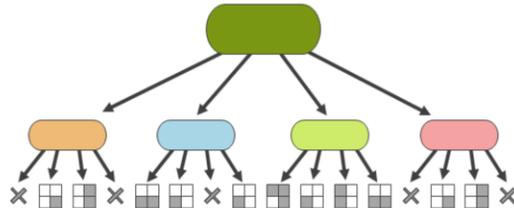
EE

Then, the first step is to expand nodes so they have eight pointers (one per possible child)

While building, we store null pointers for non-existing children, so the nodes are easily indexable

Also, we write the indexes as offsets into each *level*

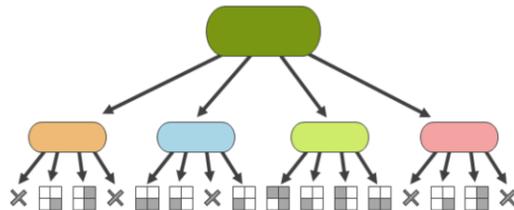
Creating a DAG from an SVO



EE

Now we create a list with the index of each node in the leaf level

Creating a DAG from an SVO



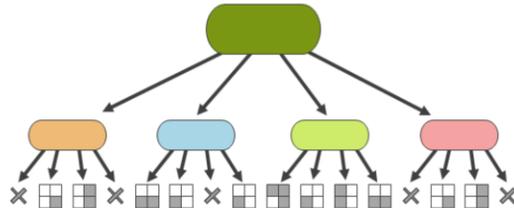
Level:

0	1111 000 001 002 003																					
1	0110	-	000 001	-	1101	002 003	-	004	1111	005 006 007 008	0110	-	009 010	-								
2	<table border="1"> <tr> <td>+</td><td>+</td><td>+</td><td>+</td><td>+</td><td>+</td><td>+</td><td>+</td><td>+</td><td>+</td><td>+</td></tr> </table>											+	+	+	+	+	+	+	+	+	+	+
+	+	+	+	+	+	+	+	+	+	+												
original idx:	0	9	1	10	2	8	3	4	6	5	7											

EE

We then sort that list, and the leaf nodes, based on the child masks in the leaf nodes. The sorted list of numbers tells us where each node was *before* sorting.

Creating a DAG from an SVO



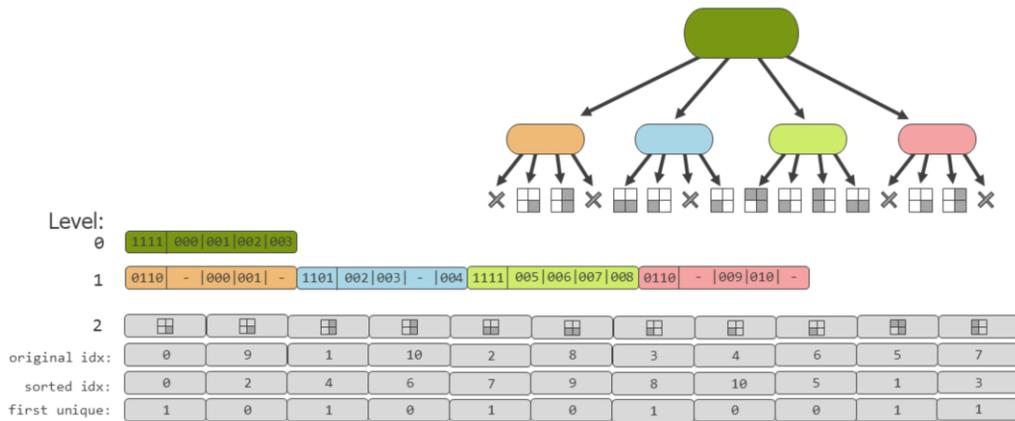
Level:

0	1111 000 001 002 003																									
1	0110	-	000 001	-	1101	002 003	-	004	1111	005 006 007 008	0110	-	009 010	-												
2	<table border="1"> <tr> <td>+</td><td>+</td><td>+</td><td>+</td><td>+</td><td>+</td><td>+</td><td>+</td><td>+</td><td>+</td><td>+</td><td>+</td><td>+</td><td>+</td><td>+</td> </tr> </table>											+	+	+	+	+	+	+	+	+	+	+	+	+	+	+
+	+	+	+	+	+	+	+	+	+	+	+	+	+	+												
original idx:	0	9	1	10	2	8	3	4	6	5	7															
sorted idx:	0	2	4	6	7	9	8	10	5	1	3															

EE

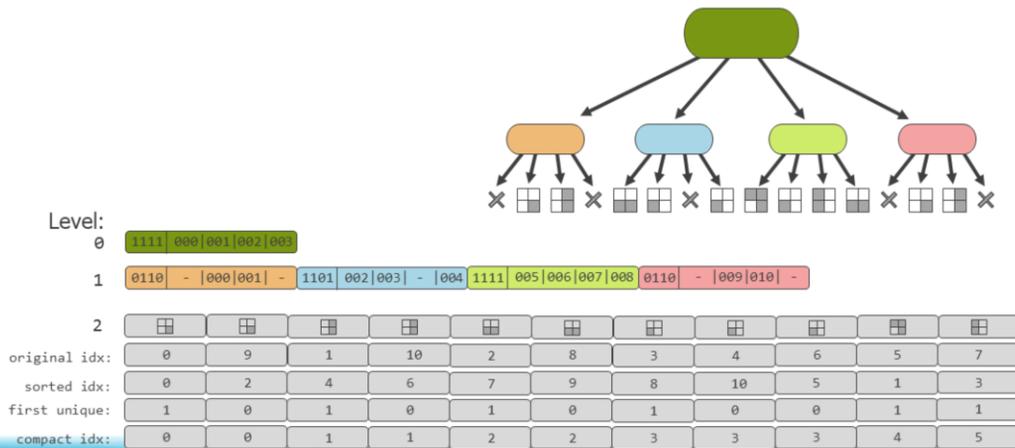
We now create the inverse of that list so we can look up where each node *ended up* after sorting

Creating a DAG from an SVO



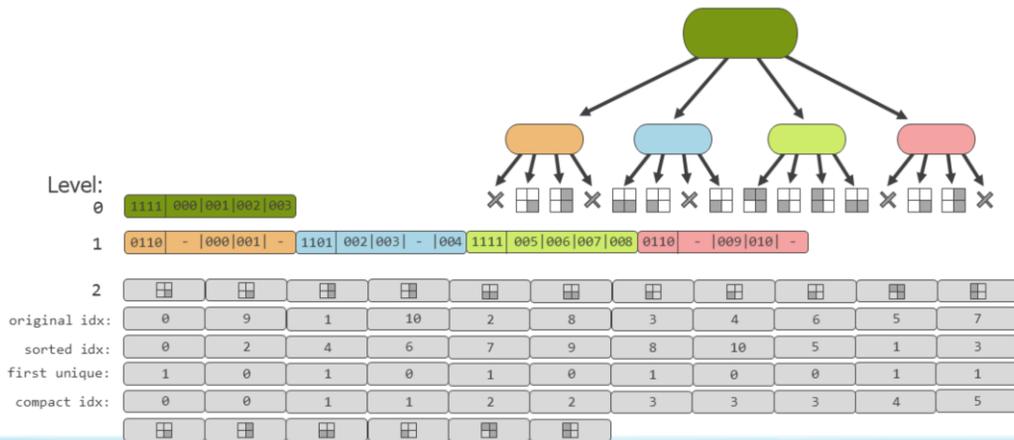
Next we run through all the leaf nodes again, and write a "1" if it is the first of its kind, and zero otherwise

Creating a DAG from an SVO



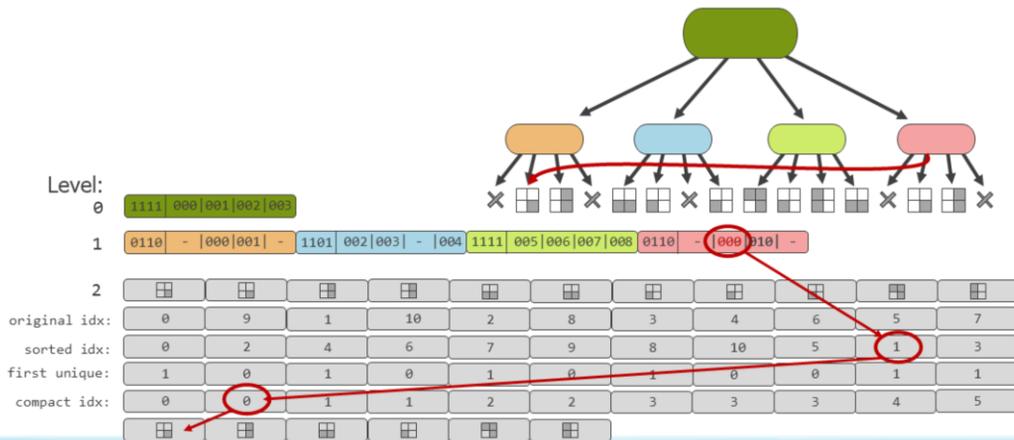
Then we calculate the prefix sum of this new array to get, for each leaf node, where the new index it should have within the leaf level

Creating a DAG from an SVO



We run a stream compaction pass on the leaf nodes, to keep only the unique nodes

Creating a DAG from an SVO



Finally, we run through all the parent pointers

<click>

We use the previous index to look up the index of the same node in the sorted list

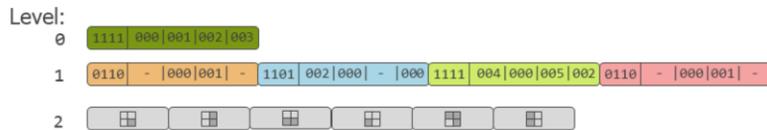
<click>

We use *that* index to find the offset in the compact list

<click>

And so we update the pointer with that index

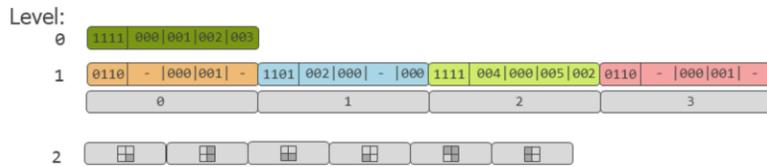
Creating a DAG from an SVO



EE

Now that we have reduced the leaf level, we can do exactly the same thing to level 1.

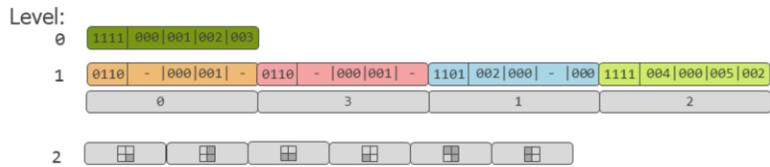
Creating a DAG from an SVO



EE

So we generate a list of indices

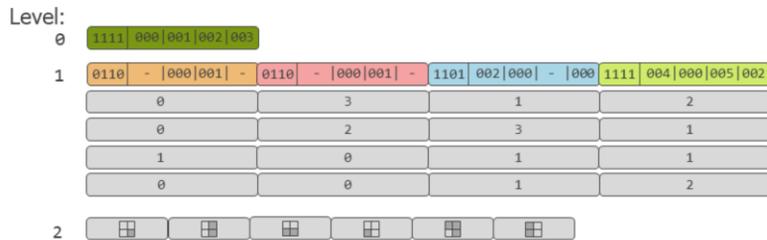
Creating a DAG from an SVO



EE

We sort the nodes. Note that this time we use ALL THE POINTERS as our sorting key.

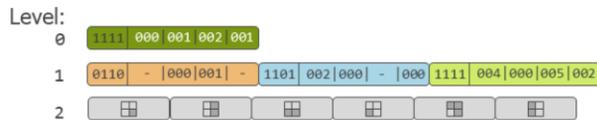
Creating a DAG from an SVO



EE

We generate the new index list, identify unique nodes and calculate the index in the compact list

Creating a DAG from an SVO



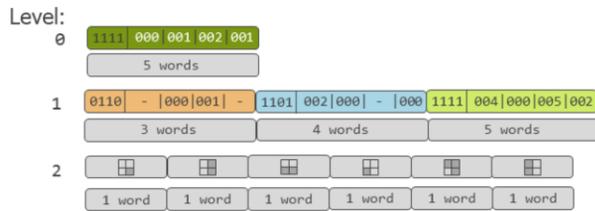
EE

Then we stream compact the nodes in level 1 and update the pointers in level 0.

In a larger SVO we would keep doing this for all levels.

When we are done, we still have a DAG with lots of NULL pointers, and the levels do not lie in a compact array.

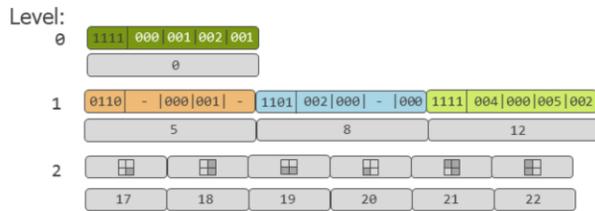
Creating a DAG from an SVO



EE

So the final step is to run through all the nodes and calculate their final size

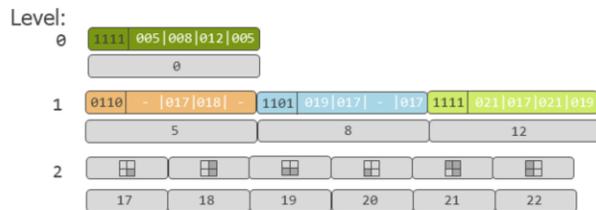
Creating a DAG from an SVO



EE

Run a prefix sum on this array to find the final index

Creating a DAG from an SVO



EE

Update the pointers to the new indices

Creating a DAG from an SVO



Issues:

- Using eight int32 numbers as a sorting key is expensive
- Since we start from the bottom of the SVO, the processing time will be dependent on the size of the leaf level

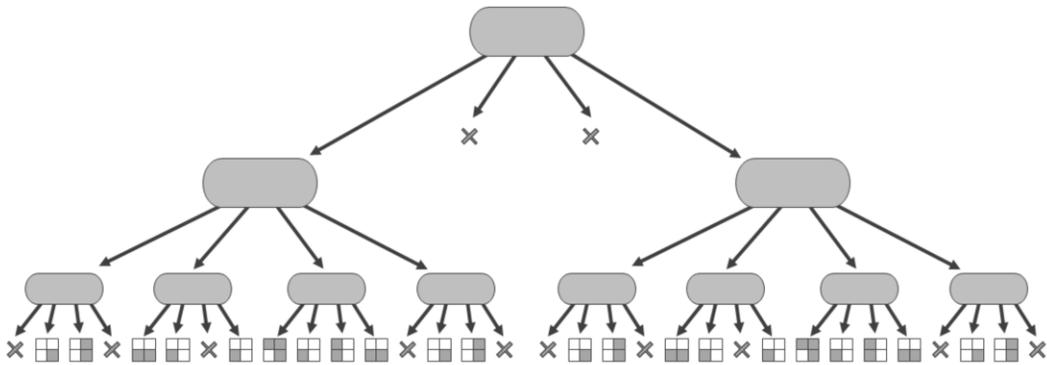
EE

And copy the nodes to their new indices in a compact array.

This algorithm works, and is really quite fast if implemented with cudpp/cub/thrust on the GPU.

But there are two annoying problems with it.

Creating a DAG from an SVO, Alternative Version



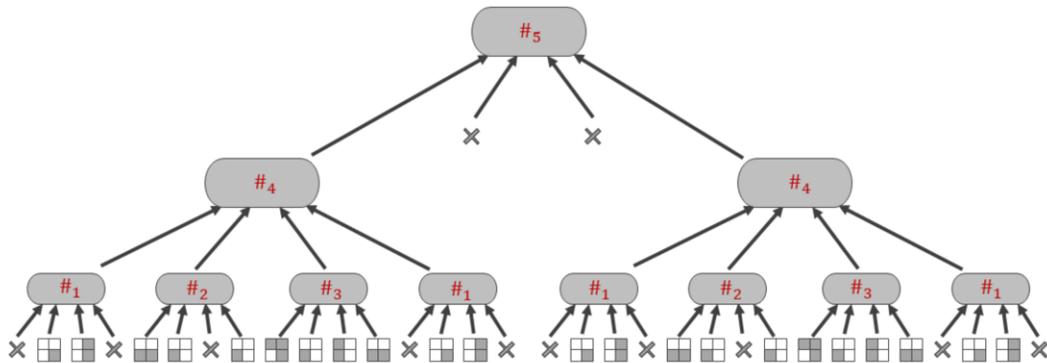
EE

While writing a TVCG extension to our I3D paper on fast construction of precomputed voxelized shadows, we came up with a much simpler and faster method.

Let's say our original SVO looks like this.

We start at level 2, and calculate a *hash value* for each node, based on the child-masks of its contained leaf nodes. <click>

Creating a DAG from an SVO, Alternative Version



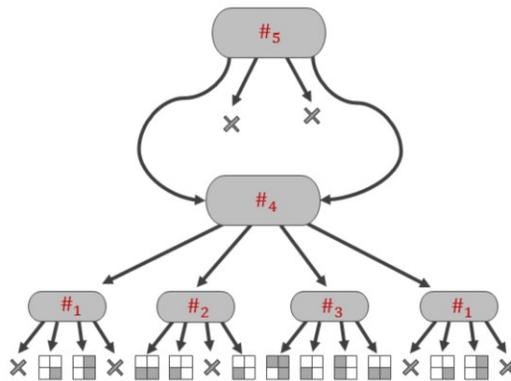
EE

We continue to create hash values for the upper levels, always creating the hash based on the hashes of the child-nodes.

Now we can remove common sub trees with a top-down algorithm instead.

Starting at level 1, we sort the nodes, based on their hash, remove redundant nodes and update the parent pointers.

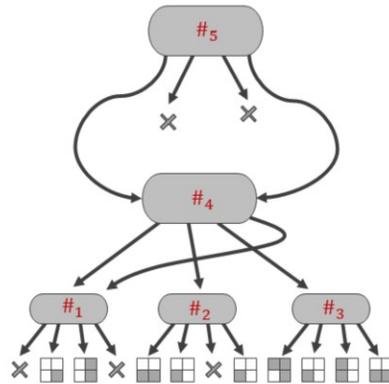
Creating a DAG from an SVO, Alternative Version



EE

Now we move on to level 2, which is a lot smaller than it was before,
sort that level, remove redundant nodes and update parent pointers...

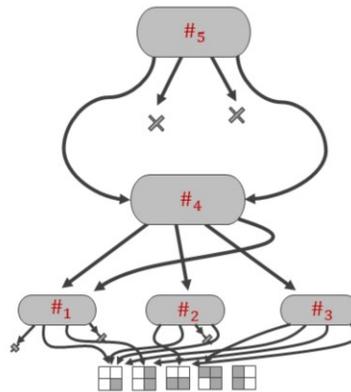
Creating a DAG from an SVO, Alternative Version



EE

That reduced the size of the leaf level, which we finally sort and reduce.

Creating a DAG from an SVO, Alternative Version



EE

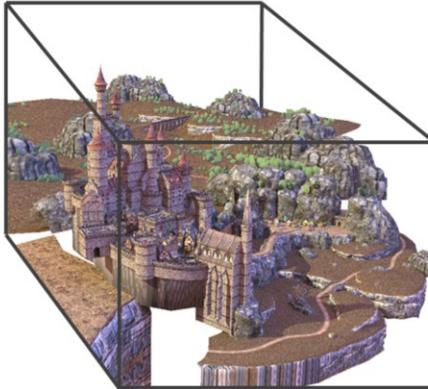
This algorithm can be much faster (was about 2x as fast in our shadow tests). Processing time will be proportional to the largest level *in the DAG* rather than in the SVO so it should scale much better. Also, we always sort on the hash value, rather than the pointers, which is much better.

There *is* a theoretical possibility that hash values could collide, which would cause very weird result, but this risk goes towards zero as we increase the size of the hash value.

For us, 64 bits was sufficient to never experience any collisions.

Compressing Large DAGs

- Epic Citadel
- 128k x 128k x 128k voxels
- DAG size ~1GB
- Pointerless (!) SVO 5.1GB

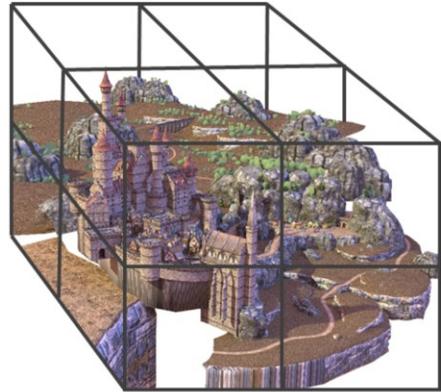
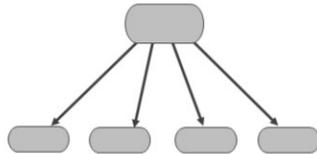


EE

At high resolutions, fitting the original SVO on GPU memory can be impossible. Since the GPU based construction algorithm also requires a fair share of temporary memory, it is often necessary to split construction into subvolumes. I will describe that briefly now.

Compressing Large DAGs

- Create a dense top of the tree:

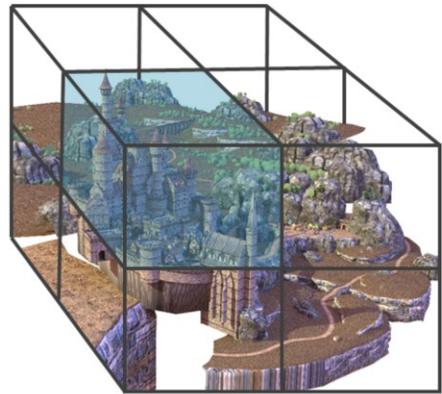
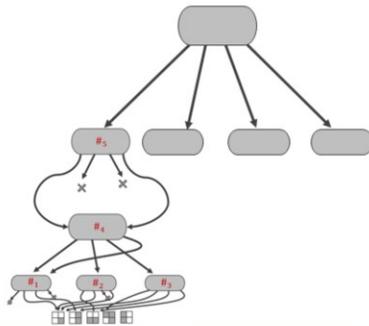


EE

At high resolutions, fitting the original SVO on GPU memory can be impossible. Since the GPU based construction algorithm also requires a fair share of temporary memory, it is often necessary to split construction into subvolumes. I will describe that briefly now.

Compressing Large DAGs

- Create a DAG

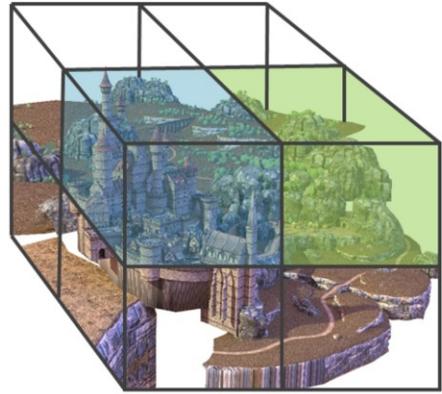
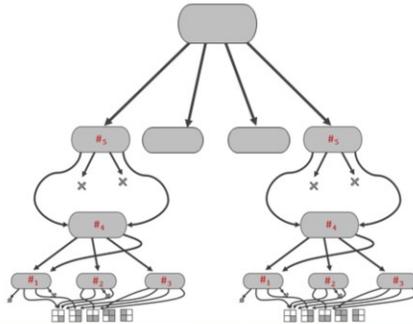


EE

At high resolutions, fitting the original SVO on GPU memory can be impossible. Since the GPU based construction algorithm also requires a fair share of temporary memory, it is often necessary to split construction into subvolumes. I will describe that briefly now.

Compressing Large DAGs

- Voxelize and Compress a second child

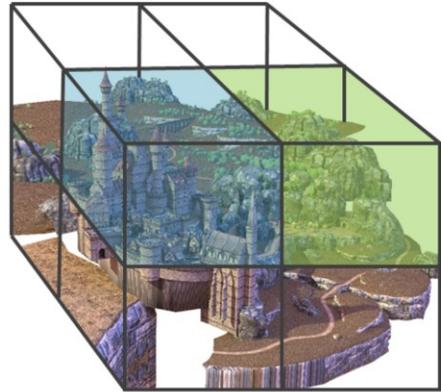
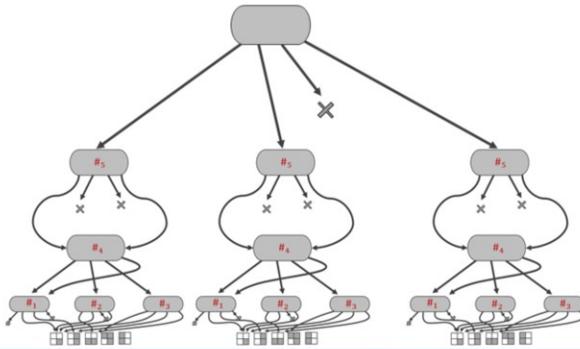


EE

At high resolutions, fitting the original SVO on GPU memory can be impossible. Since the GPU based construction algorithm also requires a fair share of temporary memory, it is often necessary to split construction into subvolumes. I will describe that briefly now.

Compressing Large DAGs

- End up with a *partially* reduced DAG



EE

When we have done this for all subvolumes, we end up with a *partially* reduced DAG. That is, all of these subtrees are optimally compressed, but there may well be subtrees in this tree that are identical to subtrees in this tree.

We already have the hash values for all the nodes, so we CAN repeat the top-down algorithm exactly as before.

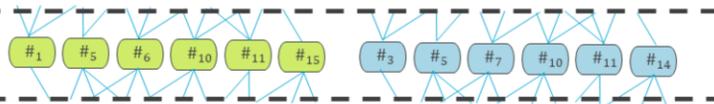
However, there is one trick that can speed things up significantly.

Merging one level

Level 2

Level 3

Level 4



EE

Since nodes in any level are already sorted, we do not have to perform a proper sort in our top down algorithm.

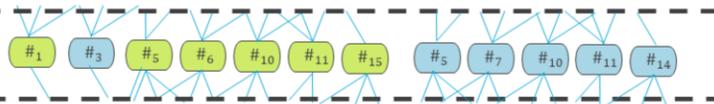
Simply merge the two levels in a linear sweep through the arrays.

Merging one level

Level 2

Level 3

Level 4



EE

Since nodes in any level are already sorted, we do not have to perform a proper sort in our top down algorithm.

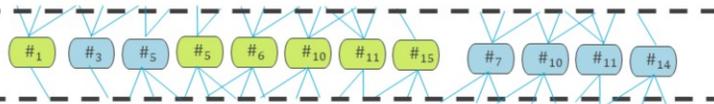
Simply merge the two levels in a linear sweep through the arrays.

Merging one level

Level 2

Level 3

Level 4



EE

Since nodes in any level are already sorted, we do not have to perform a proper sort in our top down algorithm.

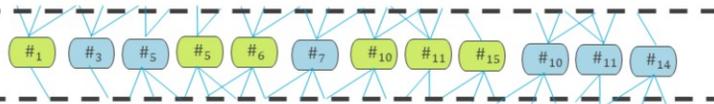
Simply merge the two levels in a linear sweep through the arrays.

Merging one level

Level 2

Level 3

Level 4



EE

Since nodes in any level are already sorted, we do not have to perform a proper sort in our top down algorithm.

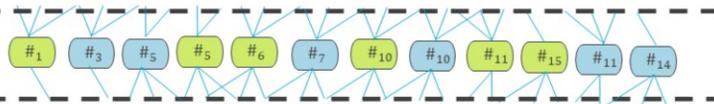
Simply merge the two levels in a linear sweep through the arrays.

Merging one level

Level 2

Level 3

Level 4



EE

Since nodes in any level are already sorted, we do not have to perform a proper sort in our top down algorithm.

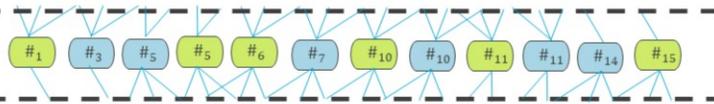
Simply merge the two levels in a linear sweep through the arrays.

Merging one level

Level 2

Level 3

Level 4



EE

Then, we simply remove redundant nodes and update the parent pointers as usual.