

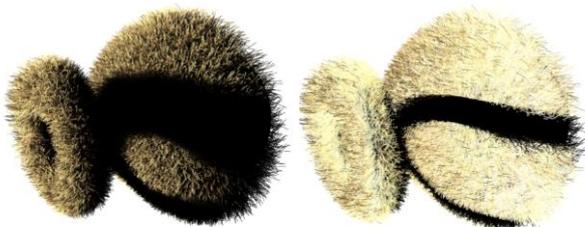
Radix sort of line primitives in CUDA for real-time Self-shadowing and Transparency in Hair

Erik Sintorn, Ulf Assarsson, Ola Olsson and Markus Billeter
Chalmers University of Technology

This poster presents an improvement of the algorithm suggested in the paper [SA08] in which the *Transform Feedback* functionality was used to sort line segments in real-time, both for the purpose of alpha blending the hair-segments (rendering back-to-front) and in order to quickly build an *opacity-map* that could be queried for self shadowing. The line segments are sorted approximately into slices and are then rendered into the corresponding slices of a 3D-texture. Here, we utilize a very fast parallel *stream-split* operation introduced in [BOA09], to quickly radix sort the line-primitives in CUDA.



Motivation

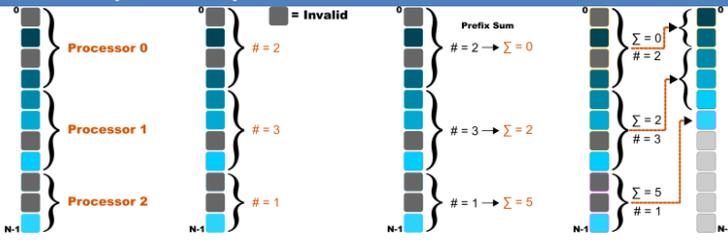


The images above illustrate the importance of self-shadowing. In the image to the right, shadows are only cast by opaque objects.



These images show hair rendered with alpha blending (left) and without (right). MSAA alone will not suffice, both because hair is really semi transparent and because the expected error of anti-aliasing is much larger for thin lines than for triangles.

Stream compaction, split and sort



Setup

Given N input elements and P warps, each warp is assigned (N/P) elements.

Count elements

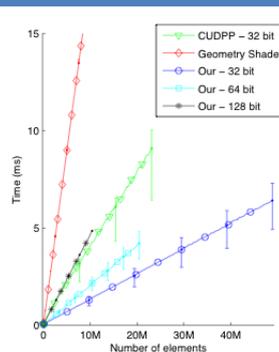
Each warp counts the number of valid elements in its assigned chunk.

Find Offsets

A prefix sum is calculated over the P values from the last step, to find output offsets.

Copy Elements

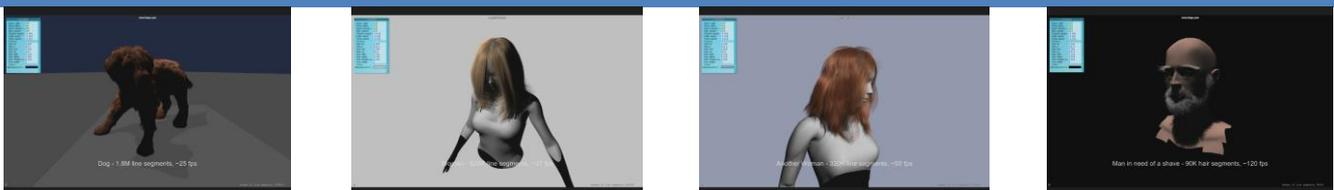
Since each warp knows the offsets for its source and destination, elements can now be copied without synchronization.



Split and sort

Given this algorithm for compacting a stream it is fairly simple to construct a *split* operation since only the final step needs to be modified to copy the invalid elements as well. By repeatedly applying this very fast operation for each bit of the integer key, we can trivially construct a radix-sort. So, to sort the lines into 256 slices, we need to split the stream eight times.

Results



Results for different models. From left to right we have (#line segments/fps): Dog (1.8M/25), Woman1(620K/37fps), Woman2(320K/50fps), Beard(90K/120fps)

References

- [SA08] Erik Sintorn, Ulf Assarsson. *Real-Time Approximate Sorting for Self Shadowing and Transparency in Hair Rendering*. Proceedings of the Symposium on Interactive 3D Graphics and Games (I3D 2008), pp 157-162, February, 2008.
- [BOA09] Markus Billeter, Ola Olsson, Ulf Assarsson. *Efficient Stream Compaction on Wide SIMD Many-Core Architectures*. High Performance Graphics 2009, August, 2009.