Cache-Oblivious Layouts of Polygonal Meshes

Sung-Eui Yoon[1], Peter Lindstrom[2], Valerio Pascucci[2], Dinesh Manocha[1]

1: University of North Carolina at Chapel Hill

2: Lawrence Livermore National Laboratory

Abstract: *We present a novel method for computing cache-oblivious layouts of meshes that improve the performances of interactive visualization and geometric processing algorithms. Our algorithm relies on a novel cache-oblivious metric for estimating the expected number of cache misses. Computing a coherent mesh layout is reduced to a combinatorial optimization problem. We have tested our algorithm into three different applications and observed 2 – 20 improvement without modifying codes of runtime applications*

Introduction:

- What is a metric to measure the expected number of cache misses of layouts?
- We measure the number of cache misses by using edge span distribution of the layout and cache miss ratio functions.
- We cast our problem as a optimization and uses local permutations to compute a layout

Methods:

Since we do not assume any cache parameters, we probabilistically decide whether a local permutation can decrease cache misses among all the possible cache parameters.

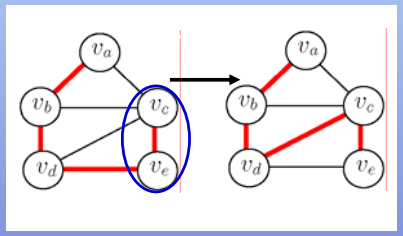

A local permutation

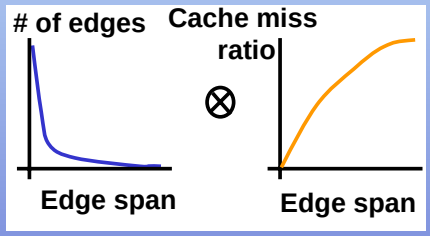

Expected number of cache misses

Results:

On the left, we compared our layout with other rendering sequences in terms of cache misses at GPU vertex cache size. On the right, performance of our layout is compared with one of depth-first layout on collision detection.

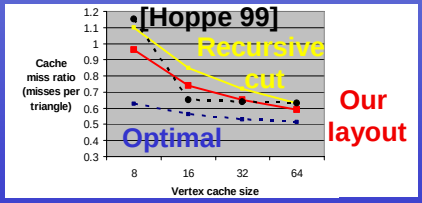

GPU vertex cache miss ratio of different layouts

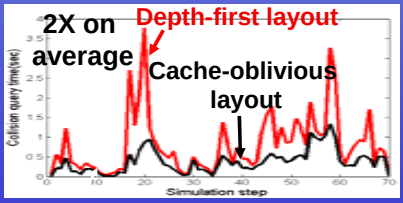

Collision detection times

Discussion:

Our layout has much shorter edge spans; therefore it reduce the number of cache misses at runtime. We would like to work on how much our layout is close to the optimal layout.

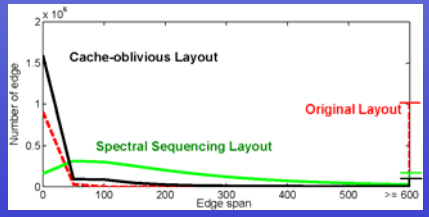

Edge span distribution of different layouts

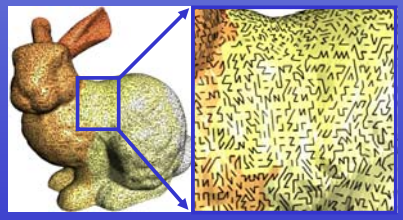

Edges having edge span 1

This work was performed under the auspices of the U.S. Department of Energy by University of California Lawrence Livermore National Laboratory under contract No. W-7405-Eng-48.

UCRL-POST-214124