## D NTNU Innovation and Creativity

## Framework for Polygonal Structures on Clusters

Leif Christian Larsen (leifchl@idi.ntnu.no) Anne C. Elster (main supervisor), Tore Fevang (co-supervisor, Schlumberger)



#### Outline

3/51

- Motivation: Pillar Gridding for Seismic Fault Detection
- Main Problem I Efficient Voxelization
- Main Problem II Voxel transfer and caching, load-balancing



4/51



(Image courtesy Schlumberger)

**D NTNU** Innovation and Creativity

5/51



#### Pillar Gridding Workflow

6/51

- 1. A geologist marks fault pillars or fault sticks
- 2. Insert additional fault pillars between interpreted pillars, and connect pillars by a polygon mesh.
- 3. Finally, polygon is *voxelized* to do computations (e.g. filtering) over voxels
- 4. Subsequently, pillars/polygon vertices can be moved



7/51

#### Problem: How can this be done when voxel data



Need efficient interface to:

Add polygons

8/51

- Remove polygons/vertices
- Move vertices
- Create/destroy meshes
- Voxelize all faces on all nodes



#### Contributions

Contributions:

- New parallel polygon voxelization algorithm
- Introduce efficient caching/transfer strategies for the operations on the previous slide
- Introduce three algorithms for load-balancing



#### 10/51

#### **Problem I: Voxelization**

Voxelization of a polygon/triangle





#### Voxelization — Introduction

- Voxelization is the process of discretizing a 3D object with 3D voxels.
- Rasterization is the same process and is implemented in hardware on GPUs, but discretizes only to 2D pixels.



#### Voxelization — Introduction

- Traditionally, visualization has focused at going the other direction (from voxels to polygons) through the Marching Cubes algorithm
- Therefore: Small research community



#### Voxelization — Introduction II

- Voxelization has applications in:
  - Certain volume-rendering algorithms in visualization
  - 3D screens (which were first suggested in 1912, but are yet to come)
  - A few other applications: seismologic interpretation for representing faults and horizons, and medical applications for representing radiation therapy beam surfaces intersecting a human body



## <sup>14/51</sup> Voxelization — Algorithms

- Most voxelization research focuses on visualization or voxelization of non-planar objects. Not relevant for our purposes
- Original voxelization algorithm by A.
   Kaufman in 1988 is one of the few voxelization techniques applicable to our problem



## <sup>15/51</sup> Kaufman's Algorithm





## <sup>16/51</sup> Voxelization — Algorithms

Voxelization is just an extension of rasterization. Thus we should be able to exploit GPU hardware to do voxelization.



## Voxelization — Algorithms

- Some previous attempts have been made on doing voxelization on GPU.
  - Previous approaches focus on visualization apps



#### <sup>18/51</sup> Voxelization — Algorithms

Idea: Transform the voxelization problem to a rasterization problem





#### <sup>19/51</sup> Voxelization — Algorithms

- 1. Translate one vertex to the origin.
- 2. Rotate vertices such that the longest edge is in z = 0 plane.
- 3. Rotate about *z*-axis such that longest edge is aligned with *x*-axis.
- 4. Rotate final vertex about *x*-axis such that it too is in the z = 0 plane.
- 5. VS Kaufman: Trade more computation for less branches



#### **Voxelization — Algorithms**

Three quaternions represent the rotations

$$(1) q_i = (\mathbf{v}, s)$$

20/51

Inverse quaternion  $q_i^{-1}$ : simply negate **v** Quaternions are combined by the formula

$$(2)q_1 * q_2 = (s_1\mathbf{v}_2 + s_2\mathbf{v}_1 + \mathbf{v}_1 \times \mathbf{v}_2, s_1s_2 - \mathbf{v}_1 \bullet \mathbf{v}_2)$$



#### Voxelization — Algorithms

Corresponding  $4 \times 4$  rotation matrix

 $\begin{bmatrix} s^{2} + v_{x}^{2} - v_{y}^{2} - v_{z}^{2} & 2v_{x}v_{y} - 2sv_{z} & 2sv_{y} + 2v_{x}v_{z} & 0\\ 2sv_{z} + 2v_{x}v_{y} & s^{2} - v_{x}^{2} + v_{y}^{2} - v_{z}^{2} & 2v_{y}v_{z} - 2sv_{x} & 0\\ 2v_{x}v_{z} - 2sv_{y} & 2sv_{x} + 2v_{y}v_{z} & s^{2} - v_{x}^{2} - v_{y}^{2} + v_{z}^{2} & 0\\ 0 & 0 & 0 & 1 \end{bmatrix}$ (3)



## <sup>22/51</sup> Voxelization Algorithm

Therefore:

- 1. Rotate initial polygon vertices. From this we obtain the combined quaternion q
- 2. Apply rasterization to the transformed vertices.
- 3. For each rasterized coordinate found, multiply by rotation matrix obtained from  $q^{-1}$ .





## <sup>24/51</sup> **Preliminary Results**

In general: GPU Speedup only when polygons are about  $10^5$ ; about 2

- Arithmetic complexity of algorithm is low
- Post-processing which must be done on CPU ruins GPU speedup.
- Using multple cores gives a slowdown on Intel systems, but not on Njord (speedup = 2, max at 8 nodes, decreases for 16/32)



# Now, let's talk about something completely different



#### **Framework Operations**

Need support for the following operations:

- Create/destroy meshes
- Add/remove polygons

26/51

- Voxelize all polygons on all nodes (Extract-All Operation)
- Move vertices of polygons



#### Framework

27/51

#### Decision: Keep polygon model global

Voxels stored on node 0  $(n_0)$ 

V

**Global polygon model**  $n_0$   $n_1$ 









## <sup>28/51</sup> Framework

Use a variant of Baumgart's winged-edge model





#### Framework

- Idea: Defer all communication to when voxelizing all polygons on all nodes.
- Keeps add/remove/move.. operations quick



#### Framework — Definitions

We use some definitions:

• Centroid  $\overline{c}$  of a triangle:

(4) 
$$\overline{c} = \frac{1}{3} \left( \sum_{i=1}^{3} v_{i,x}, \sum_{i=1}^{3} v_{i,y}, \sum_{i=1}^{3} v_{i,z} \right)$$

• Coordinate space  $C_i$  of a node: Each node istores a subset  $C_i$  of the 3D space  $N \times M \times K$ . With p = ab processes,  $|C_i| = |N|/a \times |M|/b \times |K|$ .



#### <sup>31/51</sup> Coordinate Spaces





#### Framework — Definitions

• Define the *centroid node* of a triangle T with centroid  $\overline{c}$  to be the node i such that  $\overline{c} \in C_i$ .



## 33/51 Coordinate Spaces Voxels stored on node 0 (n<sub>0</sub>) Voxels stored on node 1 (n<sub>1</sub>)



**-** x



#### Framework — Definitions

- Let *F* be a face. Then let  $V^*(F)$  be the set of coordinates in *F*'s voxelization.
- Define a function V(p) which associates which a coordinate p = (x, y, z) a user-defined, application-dependent value, typically single 32-bit float
- For each face *F*, all *V*(*p*) values for
   *p* ∈ *V*<sup>\*</sup>(*F*) must be copied/available to/on a single node in order to do computations



#### Framework — Definitions

- Define the *responsible node* of a triangle T the node at which all V(p) for  $p \in V^*(T)$  values are available after doing the Extract-All operation.
- For now, assume that the responsible node is equal to the centroid node for all triangles.



#### 36/51 **Coordinate Spaces** Voxels stored on node 0 $(n_0)$ Voxels stored on node 1 $(n_1)$ $n_0$ $n_1$ $\square$ Centroid $\overline{c}$ of triangle $C_0$ $C_1$ y Since $\overline{c} \in C_1$ , node 1 is the centroid node of the triangle **-** x



## <sup>37/51</sup> Caching Strategy I

On each node:

- 1. Voxelize all faces that node is responsible for.
- 2. Request missing V(p) values and surrounding  $k \ge 0$  V(p) values from other nodes, receive requests from other nodes.
- 3. Transfer/receive voxels to/from other nodes, caching the voxels.
- 4. Do computation.



## <sup>38/51</sup> Caching Strategy I

- $\Box$  Voxels stored on node 0 ( $n_0$ )
- Solution Voxels stored on node 1  $(n_1)$
- Voxel copied from  $n_0$  to  $n_1$  and cached on  $n_1$



NTNU, May 31st, 2007

Innovation and Creativity

39/51

#### **Caching Strategy II**

On each node:

- 1. For all faces, check if the bounding volume of the face intersects the coordinate space of the present node *and* has a different responsible node.
- 2. Transmit those voxels to the other node, if not already transmitted before.
- 3. (As nonblocking call in Step 2 proceeds) Voxelize all polygons node is responsible for.
- 4. Do computation when nonblocking call is complete.

Innovation and Creativity

## <sup>40/51</sup> Caching Strategy II

- $\Box$  Voxels stored on node 0 ( $n_0$ )
- Voxels stored on node  $1(n_1)$
- Voxels in  $B_{0,1}$  and stored on  $n_0$





#### <sup>41/51</sup> Load-Balancing

#### Question: Should the responsible node =

 $\Box$  Voxels stored on node 0 ( $n_0$ )

Solution Voxels stored on node 1  $(n_1)$ 





#### <sup>42/51</sup> Load-Balancing

- Let A(F) be the area of a set of faces F,  $\overline{A}$  = area of all faces/# of nodes,  $F^i$  = set of faces node *i* is responsible for
- Node *i* is overloaded if

(5) 
$$L(i) = A(F^i) - \overline{A}$$

is > 0. Generally use a threshold  $\delta \ge 0$  so that *i* is considered overloaded only if  $L(i) > \delta$ .



#### 43/51 Load-Balancing

- All LB algorithms execute w/ no communication.
- Simplifies, minimizes overhead, and all nodes know which node is responsible for allfaces.
- Potential disadvantage when scaling to 10,000+++ processors



#### 44/51 Load-Balancing Strategy Global

Overloaded node to which a triangle is added

The most underloaded node, to which the triangle from the **T** overloaded node is assigned

х

NTNU, May 31st, 2007

Innovation and Creativity

#### 45/51 Load-Balancing Strategy II: Local

Overloaded centroid node

x

The most underloaded node of the nodes whose NTNU coordinate space the face intersects, and which is assigned to INNOVATION and Creativity be responsible for the face

#### <sup>46/51</sup> Load-Balancing Manhattan

y

► x

#### 



NTNU, May 31st, 2007

**III**:

Strategy

#### Summary

47/51

- New parallel voxelization algorithm
- Caching/transfer strategies for the optimizing parallel voxelization operation
  - Strategy 1: Extract only required voxels (+ more if using greater block size)
  - Strategy 2: Based on bounding volumes, less data transferred + voxelization/transfer in parallel
- Load-balancing: global, local, Manhattan



#### Future Work — Voxelization

- Compare to Kaufman's algorithm.
- Extend algorithm to not only voxelize triangles but entire meshes.



#### Future Work — Framework

- Further LB investigation.
  - How often should LB be done?
  - New LB algorithms which explicitly consider how many voxels have already been cached at target node
  - Make LB algorithms fully dynamic



#### Future Work — Framework

- Investigate parameters and optimal cache/LB strategies for various:
  - data set size, average polygon size, average distance moved when moving a vertex, number of nodes, number of voxelization threads on Njord, different workloads (we consider voxelize-move-voxelize-move....), different LB thresholds



51/51

#### Thank you for your attention!

