

Parallel d -D Delaunay Triangulations in Shared and Distributed Memory

Daniel Funke and Peter Sanders

Jared Di Carlo

March 14, 2019



Triangulation in Arbitrary Dimensions

Definition

A k -simplex is a k -dimensional generalization of a triangle.





3D Triangulation



Delaunay Triangulation in Arbitrary Dimensions

Definition

In \mathbb{R}^n , given a set S of points, the *Delaunay Triangulation* of S is a triangulation of S such that the circumsphere of any simplex in $\mathcal{D}(S)$ does not contain any other point in S .




Delaunay Triangulation in Arbitrary Dimensions

- ▶ In \mathbb{R}^2 , Delaunay maximizes the minimum angle
- ▶ In \mathbb{R}^3 , it doesn't
- ▶ In \mathbb{R}^n , it minimizes the maximum *containment sphere*
- ▶ Number of tetrahedra in 3D Delaunay tetrahedralization
 - ▶ $O(n)$: uniform volume
 - ▶ $O(n \log n)$: uniform surface
 - ▶ $O(n^2)$: worst case (two lines)


Sequential 3D Delaunay

- ▶ Randomized insertion
- ▶ Flipping
- ▶ Watson Method


Watson Method




Strategy




(a) partitioning



(b) partial DTs



(c) border DT



(d) final DT

Part 1: Split

```
1: if  $n < N \vee r = \log P$  then  
2:   return sequentialDelaunay( $\mathbf{P}$ )  
3:  $k \leftarrow \text{splittingDimension}(\mathbf{P})$   
4:  $(\mathbf{P}_1 \ \mathbf{P}_2) = (p_1 \ \cdots \ p_s \mid p_{s+1} \ \cdots \ p_n) \leftarrow \text{divide}(\mathbf{P}, k)$   
5:  $\mathbf{T} = (T_1 \ T_2) \leftarrow (\text{Delaunay}(\mathbf{P}_1, r + 1) \ \text{Delaunay}(\mathbf{P}_2, r + 1))$ 
```

- ▶ Constant
- ▶ Alternating
- ▶ Largest

Part 2: Finding Border Triangles

```
6:  $\mathbf{B} \leftarrow \emptyset$ ;  $\mathbf{Q} \leftarrow \text{convexHull}(T_1) \cup \text{convexHull}(T_2)$  ▷ initialize  
7: parfor  $s_{i,x} \in \mathbf{Q}$  do ▷ si  
8:    $\text{mark}(s_{i,x})$   
9:   if  $\text{circumsphere}(s_{i,x}) \cap \text{boundingBox}(T_j) \neq \emptyset$ , with  $i \neq j$  then  
10:      $\mathbf{B} \cup= \{s_{i,x}\}$  ▷ circumsphere intersects ot  
11:     for  $s_{i,y} \in \text{neighbors}(s_{i,x}) \wedge \neg \text{marked}(s_{i,y})$  do  
12:        $\mathbf{Q} \cup= s_{i,y}$ ;  $\text{mark}(s_{i,y})$   
13:  $T_B \leftarrow \text{Delaunay}(\text{vertices}(\mathbf{B}), r + 1)$ 
```

Part 3: Merge

```
14:  $T \leftarrow (T_1 \cup T_2) \setminus \mathbf{B}; \quad \mathbf{Q} \leftarrow \emptyset$   
15: parfor  $s_b \in T_B$  do  
16:   if  $\text{vertices}(s_b) \not\subset \mathbf{P}_1 \wedge \text{vertices}(s_b) \not\subset \mathbf{P}_2$  then  
17:      $T \cup= \{s_b\}; \quad \mathbf{Q} \cup= \{s_b\}$   
18:   else  
19:     if  $\exists s \in \mathbf{B} : \text{vertices}(s) = \text{vertices}(s_b)$  then  
20:        $T \cup= \{s_b\}; \quad \mathbf{Q} \cup= \{s_b\}$ 
```

Part 4: Update Data Structures

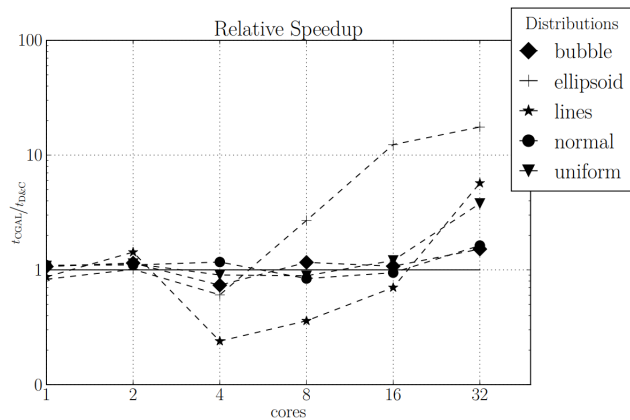
Neighborhood update:

```
21: parfor  $s_x \in Q$  do
22:   for  $d \in \{1, \dots, D + 1\}$  do
23:     if  $\text{neighbors}_d(s_x) \notin T$  then
24:        $C \leftarrow \{s_c : f_d(s_x) = f_d(s_c)\}$ 
25:       for  $s_c \in C$  do
26:         if  $|\text{vertices}(s_x) \cap \text{vertices}(s_c)| = D$  then
27:            $\text{neighbors}_d(s_x) \leftarrow s_c; \quad Q \cup = s_c$ 
28: return  $T$ 
```


Data Structures

- ▶ Check for duplicates
 - ▶ Hash table of simplices
- ▶ Update Neighbors
 - ▶ Face hash table
 - ▶ Build during Part 2
- ▶ Merge Large Blocks
 - ▶ Binary Search Tree
 - ▶ $O(\log k)$ access after k merges

Results



Results



Comparison Of Results

- ▶ Tested on 1 million uniformly distributed points (same dataset)
- ▶ Different output size depending on # of threads?
- ▶ Volume is wrong
- ▶ Intel E5-2699V3 (2014, 2.30 GHz, 18 cores), 32 GB RAM (nobody used this much)
- ▶ Very hard to find program (author has deleted published GitHub page)

	1 Core	16 Core
This Paper	233 s	42 s
CGAL	8 s	n/a
Tetgen	5 s	n/a

Questions

- ▶ Results?
 - ▶ Paper: 50,000,000 points in 64 seconds on 32 cores (24k pts/(sec cpu))
 - ▶ My result: 1,000,000 points in 42 seconds on 16 cores (1.4k pts/(sec cpu))
- ▶ Do we actually need Delaunay?
 - ▶ Delaunay algorithms are $O(n \log n)$
 - ▶ Can get a “good” tetredralization in $O(n)$ time...